arXiv:1606.00195v2 [cs.DC] 6 Dec 2016

Self-Stabilizing Reconfiguration
(Technical Report)

Shlomi Dolev * Chryssis Georgiou † Ioannis Marcoullis† Elad M. Schiller ‡

Abstract

Current reconfiguration techniques are based on starting the system in a consistent configuration, in which all participating entities are in a predefined state. Starting from that state, the system must preserve consistency as long as a predefined churn rate of processors joins and leaves is not violated, and unbounded storage is available. Many working systems cannot control this churn rate and do not have access to unbounded storage. System designers that neglect the outcome of violating the above assumptions may doom the system to exhibit illegal behaviors. We present the first automatically recovering reconfiguration scheme that recovers from transient faults, such as temporal violations of the above assumptions. Our self-stabilizing solutions regain safety automatically by assuming temporal access to reliable failure detectors. Once safety is re-established, the failure detector reliability is no longer needed. Still, liveness is conditioned by the failure detector's unreliable signals. We show that our self-stabilizing reconfiguration techniques can serve as the basis for the implementation of several dynamic services over message passing systems. Examples include self-stabilizing reconfigurable virtual synchrony, which, in turn, can be used for implementing a self-stabilizing reconfigurable state-machine replication and self-stabilizing reconfigurable emulation of shared memory.

Keywords: Self-stabilization, Dynamic Participation, Reconfiguration, Virtual Synchrony, State Machine Replication.

*Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva, Israel. Email: dolev@cs.bgu.ac.il

†Department of Computer Science, University of Cyprus, Nicosia, Cyprus. Email: {chryssis, imarco01}@cs.ucy.ac.cy . Supported by the University of Cyprus.

‡Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, SE-412 96, Sweden. Email: elad@chalmers.se.

1 Introduction

Motivation. We consider distributed systems that work in dynamic asynchronous environments, such as a shared storage system [20]. Quorum configurations [21, 23], i.e., set of active processors (servers or replicas), are typically used to provide services to the participants of the system. Since over time, the (quorum) configuration may gradually lose active participants due to voluntary leaves and stop failures, there is a need to allow the participation of newly arrived processors and from time to time to *reconfigure* so that the new configuration is built on a more recent participation group. Over the last years, a number of reconfiguration techniques have been proposed, mainly for state machine replication and emulation of atomic memory (e.g., [1–5, 7, 8, 16–19, 22]). These reconfiguration techniques are based on starting the system in a consistent configuration, in which all processors are in their initial state. Starting from that state, the system must preserve consistency as long as a predefined churn rate of processors' joins and leaves is not violated and unbounded storage is available. Furthermore, they do not tolerate *transient faults* that can cause an arbitrary corruption of the system's state.

Many working systems cannot control their churn rate and do not have access to unbounded storage. System designers that neglect the outcome of violating the above assumptions may doom the system to forever exhibit a behavior that does not satisfy the system requirements. Furthermore, the dynamic and difficult-to-predict nature of distributed systems gives rise to many fault-tolerance issues and requires efficient solutions. Large-scale message passing networks are asynchronous and they are subject to transient faults due to hardware or software temporal malfunctions, short-lived violations of the assumed failure rates or violation of correctness invariants, such as the uniform agreement among all current participants about the current configuration. Fault tolerant systems that are *self-stabilizing* [9] can recover after the occurrence of transient faults (as long as the program's code is still intact).

Our contributions and approach. We present the first automatically recovering reconfiguration scheme that recovers from transient faults, such as temporary violations of the predefined churn rate or the unexpected activities of processors and communication channels. Our blueprint for self-stabilizing reconfigurable distributed systems can withstand a temporal violation of such assumptions, and recover once conditions are resumed. It achieves this with only a bounded amount of local storage and message size. Our self-stabilizing solutions regain safety automatically by assuming temporal access to reliable failure detectors. Once safety is re-established, the failure detector reliability is no longer needed; still, liveness is conditioned by the failure detector's unreliable signals. We now overview our approach.

Reconfiguration scheme: Our scheme comprises of two layers that appear as a single "black-box" module to an application that uses the reconfiguration service. The objective is to provide to the application a *conflict-free* configuration, such that no two alive processors consider different configurations. The first layer, called *Reconfiguration Stability Assurance* or *recSA* for short (detailed in Section 3.1), is mainly responsible for detecting configuration conflicts (that could be a result of transient faults). It deploys a *brute-force* technique for converging to a conflict-free new configuration. It also employs another technique for *delicate* configuration replacement when a processor notifies that it wishes to replace the current configuration with a new set of participants. For both techniques, processors use a failure detector (detailed in Section 2) to obtain membership information, and configuration convergence is reached when failure detectors have temporal reliability. Once a uniform configuration is installed, the failure detectors' reliability is no longer needed and from then on our liveness conditions consider unreliable failure detectors. The decision for requesting a delicate reconfiguration is controlled by the other layer, called *Reconfiguration Management* or *recMA* for short (detailed in Section 3.2).

Specifically, if a processor suspects that the dependability of the current configuration is under jeopardy, it seeks to obtain a majority approval from the alive *members* of the current configuration, and request a (delicate) reconfiguration from *recSA*. Moreover, in the absence of such a majority (e.g., configuration replacement was not activated "on time" or the churn assumptions were violated), the *recMA* can aim to control the recovery via an *recSA* reconfiguration request. Note that the current participant set can, over time, become different than the configuration member set. As new members arrive and others leave, changing the configuration based on system membership would imply a high frequency of (delicate) reconfigurations, especially in the presence of high churn. We avoid unnecessary reconfiguration requests by requiring a weak liveness condition: if a majority of the configuration set has not collapsed, then there exists at least one processor that is known to trust this majority in the failure detector of each alive processor. Such active configuration members can aim to replace the current configuration with a newer one (that would provide an approving majority for prospective reconfigurations) without the use of the brute-force stabilization technique.

Joining mechanism: We complement our reconfiguration scheme with a self-stabilizing joining mechanism (detailed in Section 3.3) that manages and controls the inclusion of new processors into the system. Here extra care needs to be taken so that newly joining processors do not "contaminate" the system state with stale information (due to arbitrary faults). For this, together with other techniques, we follow a snap-stabilizing data link protocol (see Section 2). We have designed our joining mechanism so that the decision of whether new members should be included in the system or not is *application-controlled*. In this way, the churn (regarding new arrivals) can be "fine-tuned" based on the application requirements; we have modeled this by having joining processors obtaining approval from a majority of the members of the current configuration (if no reconfiguration is taking place). These, in turn, provide such approval if the application's (among other) criteria are met. We note that in the event of transient faults, such as an unavailable approving majority, *recSA* ensures recovery via brute-force stabilization that includes all alive processors.

Applications: We demonstrate the usability and modularity of our self-stabilizing reconfiguration scheme and joining mechanism by using them to develop self-stabilizing dynamic participation versions of several algorithms: a label algorithm for providing a bounded self-stabilizing labeling scheme (Section 4.1); a self-stabilizing counter increment algorithm (Section 4.2); a self-stabilizing virtual synchrony algorithm that leads to self-stabilizing state machine replication and a self-stabilizing MWMR emulation of shared memory (Section 4.3). These algorithms are derived by combining our reconfiguration scheme and joining mechanism with the corresponding self-stabilizing algorithms developed for static membership systems in [11].

Related work. As mentioned, existing solutions for providing reconfiguration in dynamic systems, such as [17] and [2], do not consider transient faults and self-stabilization, because their correctness proofs (implicitly) depend on a coherent start [20] and also assume that fail-stops can never prevent the (quorum) configuration to facilitate configuration updates. They also often use unbounded counters for ordering consensus messages (or for shared memory emulation) and by that facilitate configuration updates, e.g., [17]. Our self-stabilizing solution recovers after the occurrence of transient faults, which we model as an arbitrary starting state, and guarantees a consistent configuration that provides (quorum) services, e.g., allowing reading from and writing to distributed shared memory objects, and at the same time managing the configuration providing these services.

Significant amount of research was dedicated in characterizing the fault-tolerance guarantees that can be provided by difference quorum system designs; see [23] for an in depth discussion. In this paper we use majorities, which is regarded as the simplest form of a quorum system (each set composed of a majority of the processors is a quorum). Our reconfiguration scheme can be modified to support more complex, quorum

systems, as long as processors have access to a mechanism (a function actually) that given a set of processors can generate the specific quorum system. Another important design decision is *when* a reconfiguration (delicate in our case) must take place; see the related discussion in [20]. One simple decision would be to reconfigure when a fraction (e.g., 1/4th) of the members of a configuration appear to have failed. More complex decisions could use prediction mechanisms (possibly based on statistics). This issue is outside of the scope of this work; however, we have designed our reconfiguration scheme (specifically the *recMA* layer) to be able to use any decision mechanism imposed by the application (via an application interface).

2 System Settings

Processing entities. We consider an asynchronous message-passing system of processors. Each processor p_i has a unique identifier, i, taken from a totally-ordered set of identifiers P. The number of live and connected processors at any time of the computation is bounded by some N such that $N \ll |P|$. We refer to such processors as *active*. We assume that the processors have knowledge of the upper bound N, but not of the actual number of active processors. Processors may stop-fail by crashing; a processor may crash at any point without warning, and in this event a crashed processor takes no further steps and never rejoins the computation. (For readability sake, we model rejoins as transient faults rather than considering them explicitly. Self-stabilization inherently deals with rejoins by regarding the past join information as possibly corrupted.) New processors may join the system (using a joining procedure) at any point in time with an identifier drawn from P, such that this identifier is only used by this processor forever. A *participant* is an active processor that has joined the computation. Note that N accounts for all active processors, both participants and those that are still joining.

Communication. The network topology is that of a fully connected graph, and links have a bounded capacity cap. Processors exchange low-level messages called *packets* to enable a reliable delivery of high level *messages*. Packets sent may be lost, reordered, or duplicated but not arbitrarily created, although the channels may initially (after transient faults) contain stale packets, which due to the boundedness of the channels are also bounded in a number that is in $O(N^2 cap)$. We assume the availability of self-stabilizing protocols for reliable FIFO end-to-end message delivery over unreliable channels with bounded capacity, such as the ones of [10] or [12].

Specifically, when processor p_i sends a packet, pkt_1, to processor p_j, the operation $send$ inserts a copy of pkt_1 into the FIFO queue representing the communication channel from p_i to p_j. Since links are bounded in capacity, the new packet might be omitted or some already sent packet may be lost. While we assume that packets can spontaneously be omitted, i.e., lost from the channel, a packet that is sent infinitely often is received infinitely often. Namely, the communication channels provide *fair communication*. The policy of acknowledging is that acknowledgments are sent only when a packet arrives, and not spontaneously. Packet pkt_1 is retransmitted until more than the total capacity acknowledgments arrive, and then pkt_2 starts being transmitted. This forms an abstraction of token carrying messages between the two processors. In this way the two processors (sender and receiver) can continuously exchange a "token". We use this token exchange technique to implement a *heartbeat* for detecting whether a processor is active or not; when a processor in no longer active, the token will not be returned back to the other processor.

Due to the possibility of arbitrary faults and of the dynamic nature of the network, we cannot assume that processors have knowledge of the identifier of the processor with which they are communicating. We employ two anti-parallel data-link protocols, where every packet of one data-link is identified by the identifiers of the sender and receiver of the data link it participates in. For example, if the communication link connects

p_i and p_j, packets of the data link in which p_i (p_j) acts as the sender that traverse from p_i to p_j (p_j to p_i) are identified by the label p_i (p_j), while the label of packets traversing from p_j (p_i) are extended by adding p_j (p_i) to the label to form the label p_x, p_j (p_x, p_i, respectively). Any packet p_x, p_y arriving to p_i (p_j) where $x \neq i$ ($x \neq j$) is ignored. Thus, eventually the data link in which p_i is the sender is implemented by packets with label p_i (p_j) traversing from p_i to p_j (p_j to p_i). The analogous holds for the packets implementing the data link in which p_j serves as the sender. Thus, both parties will eventually know the identifier of the other party and regard the token of the data link in which the sender has the greater identifier among them, to be the used token.

Using the underlying packet exchange protocol described, a processor p_i that has received a packet from some processor p_j which did not belong to p_i's failure detector, engages in a two phase protocol with p_j in order to "clean" their intermediate link. This is done before any messages are delivered to the algorithms that handle reconfiguration, joining and applications. We follow the snap-stabilizing data link protocol detailed in [15]. A *snap-stabilizing* protocol is one which allows the system (after faults cease) to behave according to its specification upon its first invocation. We require that every data-link established between two processors is initialized and cleaned straight after it is established. In contrast to [15] where the protocol is run on a tree and initiated from the root, our case requires that each pair of processors takes the responsibility of cleaning their intermediate link. Snap-stabilizing data links do not ignore signals indicating the existence of new connections, possibly some physical carrier signal from the port. In fact, when such a connection signal is received by the newly connected parties, they start a communication procedure that uses the bound on the packet in transit and possibly in buffers too, to clean all unknown packets in transit, by repeatedly sending the same packet until more than the round trip capacity acknowledgments arrive.

(N, Θ)-failure detector. We consider the (N, Θ)-failure detector that uses the token exchange and heartbeat detailed above. This is an extension of the Θ-failure detector used in [6]. It allows each processor p_i to order other processors according to how recently they have communicated. Each processor p_i maintains an ordered heartbeat count vector $nonCrashed$, with an entry corresponding to each processor p_k that exchanges the token (i.e., sends a heartbeat) with p_i. Specifically, whenever p_i receives the token from p_j, it sets the count corresponding to p_j to 0 and increments the count of every other processor by one. In this way, p_i manages to rank every processor p_k according to the token exchanges that it has performed with p_i in relation to the token exchanges that it has performed with some other processor p_j. So the processor that has most recently contacted p_i is the first in p_i's vector.

The technique enables p_i to obtain an estimate on the number of processors n_i that are active in the system; $n_i \leq N$. Assuming that p_c is the most recently crashed processor, then every processor other than p_c will eventually exchange the token with p_i many times, and their heartbeat count will be set to zero, while p_c's will be increasing continuously. Eventually, every other processor's count (given they remain alive and communicating) will become lower than p_c's and p_c will be ranked last in $nonCrashed$. Moreover, while difference between heartbeat counts of non-crashed processors does not become large, the difference of these counts and that of p_c increases to form a significant ever-expanding "gap". The last processor before the gap is the ${n_i}^{th}$ processor and this provides an estimate on the number of active processors. Since there are at most N processors in the computation at any given time, we can ignore any processors that rank below the N^{th} vector entry. If, for example, the first 30 processors in the vector have corresponding counters of up to 30, then the 31^{st} will have a count much greater than that, say 100; so n_i will be estimated at 30. This estimation mechanism is suggested in [13] and in [14].

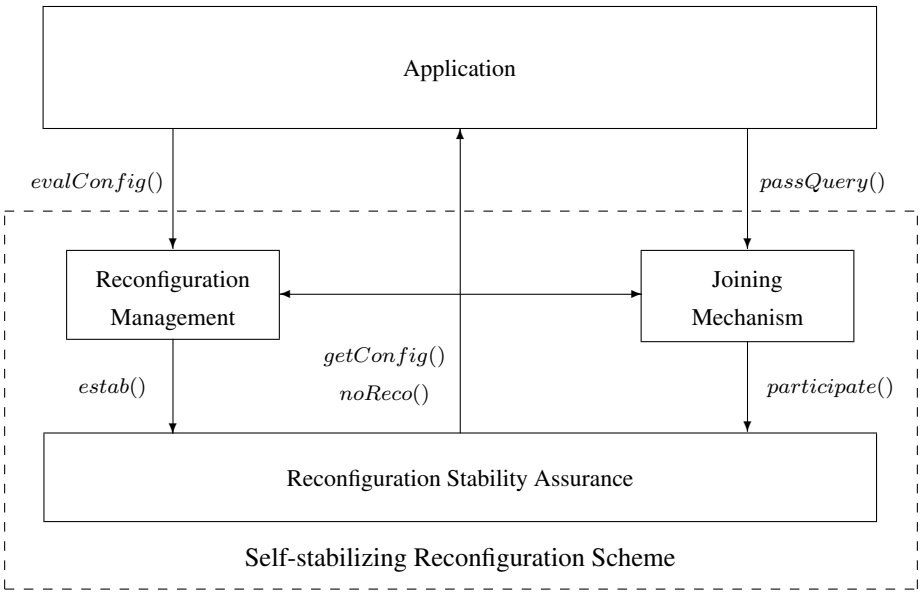

Figure 1: The reconfiguration scheme modules internal interaction and the interaction with the application. The Reconfiguration Stability Assurance ($recSA$) layer provides information on the current configuration and on whether a reconfiguration is not taking place using the $getConfig()$ and $noReco()$ interfaces. This is based of local information. The Reconfiguration Management ($recMA$) layer uses the prediction mechanism $evalConfig()$ which is application based to evaluate whether a reconfiguration is required. If a reconfiguration is required, $recMA$ initiates it with $estab()$. Joining only proceeds if a configuration is in place and if no reconfiguration is taking place. When the joining mechanism has received a permission to access the application (using $passQuery()$) it can then join via $participate()$. The direction of an arrow from a module A to a module B illustrates the transfer of the specific information from A to B.

The interleaving model and self-stabilization. A program is a sequence of *(atomic) steps*. Each atomic step starts with local computations and ends with a communication operation, i.e., packet $send$ or $receive$. We assume the standard interleaving model where at most one step is executed in every given moment. An input event can either be the arrival of a packet or a periodic timer triggering p_i to (re)send. Note that the system is asynchronous and the rate of the timer is totally unknown. The *state*, c_i, consists of p_i's variable values and the content of p_i's incoming communication channels. A step executed by p_i can change the state of p_i. The tuple of the form $(c_1, c_2, \cdots, c_n)$ is used to denote the *system state*. An *execution (or run)* $R = c_0, a_0, c_1, a_1, \ldots$ is an alternating sequence of system states c_x and steps a_x, such that each c_{x+1}, except the initial system state c_0, is obtained from c_x by the execution of a_x. A practically infinite execution is an execution with many steps, where many is defined to be proportional to the time it takes to execute a step and the life-span time of a system. The system's task is a set of executions called *legal executions* (LE) in which the task's requirements hold. An algorithm is *self-stabilizing* with respect to LE when every (unbounded) execution of the algorithm has a suffix that is in LE.

3 Self-stabilizing Reconfiguration Scheme

The reconfiguration scheme is composed of the Reconfiguration Stability Assurance ($recSA$) layer (Section 3.1), the Reconfiguration Management ($recMA$) layer (Section 3.2), and is accompanied by the Joining Mechanism (Section 3.3). Figure 1 depicts the interaction between the modules and with the application.

The Reconfiguration Stability Assurance ($recSA$) layer ensures that participants eventually have a common configuration set. It also introduces processors that want to join the computation and provides information on the current configuration and on whether a reconfiguration is not taking place using the $getConfig()$ and $noReco()$ interfaces, respectively.

The Reconfiguration Management ($recMA$) layer strives to maintain a majority of active processors of the configuration set, to this end, and may also request a reconfiguration from $recSA$ via the $estab()$ interface. This is done when a configuration majority is suspected as collapsed or if a majority of active processor configuration members appears to require a reconfiguration based on some application-defined prediction function ($evalConf()$). A joining mechanism gives the application the leverage required to control participation and ensure that processors enter the computation with the most recent state. A joiner becomes a participant via $participate()$ only if $passQuery()$ of a majority of configuration members is reported as True. We now proceed with the details.

3.1 Reconfiguration Stability Assurance

We present the Reconfiguration Stability Assurance layer ($recSA$), a self-stabilizing algorithm for assuring correct configuration while allowing the updates from the Reconfiguration Management layer (Section 3.2). We first describe the algorithm (Section 3.1.1) and then we prove its correctness (Section 3.1.2).

3.1.1 Algorithm Description

We first present an overview of the algorithm and then proceed to a line-by-line description.

Overview

The $recSA$ layer uses a self-stabilizing algorithm (Algorithm 3.1) for assuring correct configuration while allowing the updates from the reconfiguration management layer. Algorithm 3.1 guarantees that (1) all active processors have eventually identical copies of a single configuration, (2) when participants notify the system that they wish to replace the current configuration with another, the algorithm selects one proposal and replaces the current configuration with it, and (3) joining processors can become participants eventually.

The algorithm combines two techniques: one for *brute force stabilization* that recovers from stale information and a complementary technique for *delicate (configuration) replacement*, where participants jointly select a single new configuration that replaces the current one. As long as a given processor is not aware of ongoing configuration replacements, Algorithm 3.1 merely monitors the system for stale information, e.g., it makes sure that all participants have a single (non-empty) configuration. During these periods the algorithm allows the invocation of configuration replacement processes (via the $estab(\text{set})$ interface, triggered by the Reconfiguration Management layer) as well as the acceptance of joining processors as participants (via the $participate()$ interface, triggered by the Joining layer). During the process of configuration replacement, the algorithm selects a single configuration proposal and replaces the current one with that proposal before returning to monitor for configuration disagreements.

While the system reconfigures, there is no immediate need to allow joining processors to become participants. By temporarily disabling this functionality, the algorithm can focus on completing the configuration replacement using the current participant set. To that end, only participants broadcast their states at the end of the do forever loop (line 29), and only their messages arrive to the other active processors (line 30). Joining processors receive such messages, but cannot broadcast before their safe entry to the participants' set via the function $participate()$ (line 14), which enables p_i's broadcasting. Note that non-participants

Algorithm 3.1: Self-stabilizing Reconfiguration Stability Assurance; code for processor p_i

1 **Variables**: The following arrays consider both p_i's own value (entry i) and p_j's most recently received value (entry j).
2 config[]: an array in which config[i] is p_i's view on the current configuration. Note that p_i assigns the *empty (configuration)* value $\perp$ after receiving a conflicting (different) non-empty configuration value.
3 FD[]: an array in which FD[i] represents p_i's failure detector. Note that we consider only the trusted processors rather than the suspected ones. Namely, crashed processors are eventually suspected.
4 FD[]$.part$ is the participant set, where FD[i]$.part$ is an alias for $\{p_j \in \text{FD}[i] : \text{config}[j] \neq \sharp\}$ and FD[j]$.part$ refers to the last value received from p_j. Namely, the FD field of every message encodes also this participation information.
5 prp[] is an array of pairs $\langle phase \in \{0,1,2\}, set \subseteq P\rangle$, where prp[$i$] refers to p_i's configuration replacement notifications. In the pair prp[i], the field set can either be $\perp$ ('no value') or the proposed processor set.
6 echo[] is an array in which echo[i] is (FD[i]$.part$, prp[i], all[i])'s alias and echo[j] refers to the most recent value that p_i received from p_j after p_j had responded to p_i with the most recent values it got from p_i.
7 all[] is an array of Booleans, where all[i] refers to the case in which p_i observes that all trusted processors had noticed its current (maximal) notification and they hold the same notification.
8 allSeen: a list of processors p_k for which p_i received the all[k] indication.
9 **Interface functions**:
10 **function** $chsConfig() =$ **return(choose**($\{\text{config}[k]\}_{p_k \in \text{FD}[i]} \setminus \{\sharp\}$)), where **choose**($\emptyset$) $= \perp$ else **choose**(set) $\in set$;
11 **function** $getConfig() = \{$**if** $noReco()$ **then return**($chsConfig()$) **else return**(config[i])$\}$;
12 **function** $noReco() = ((p_i \notin (\cap_{p_j \in \text{FD}[i] \setminus \{p_i\}} \text{FD}[i])) \vee (|\{\text{config}[k]\}_{p_k \in \text{FD}[i]} \setminus \{\sharp\}| > 1) \vee (\{\text{FD}[i].part\} \neq \{\text{FD}[k].part, \text{echo}[k].part\}_{p_k \in \text{FD}[i]}) \vee (\nexists p_k \in \text{FD}[i] : config[k] = \perp) \vee (prp[k] \neq dfltNtf))$ /* invariant tests */
13 **function** $estab(\text{set}) = \{$**if** $(noReco() \wedge (set \notin \{\text{config}[i], \emptyset\}))$ **then** prp[i] $\leftarrow \langle 1, \text{set}\rangle\}$
14 **function** $participate() = \{$**if** $(noReco())$ **then** config[i] $\leftarrow chsConfig()\}$
15 **Constants and macros**: $dfltNtf = \langle 0, \perp\rangle$ /* the default notification tuple */
16 **macro** $degree(k) = (2 \cdot \text{prp}[k].phase + |\{1 : \text{all}[k]\}|)$ /* p_k's most-recently-received prp degree */
17 **macro** $echoNoAll(k) =$ **return** $(\{(\text{FD}[i].part, \text{prp}[i])\} = \{(\text{echo}[j].part, \text{echo}[j].\text{prp})\}_{p_j \in \text{FD}[i].part})$
18 **macro** $echo() =$ **return** $(\{(\text{FD}[i].part, \text{prp}[i], \text{all}[i])\} = \{\text{echo}[j]\}_{p_j \in \text{FD}[i].part})$
19 **macro** $same(k) =$ **return** $(\{(\text{FD}[i].part, \text{prp}[i])\} = \{(\text{FD}[k].part, \text{prp}[k])\})$
20 **macro** $maxNtf() = \{$**if** $\{\text{prp}[k]\}_{p_k \in FD[i].part} = \{dfltNtf\}$ **then return** $\perp$ **else return** $\max_{\text{lex}}\{\text{prp}[k]\}_{p_k \in FD[i].part}\}$
21 **macro** $configSet(val) = \{$**foreach** p_k **do** $(\text{config}[k], \text{prp}[k]) \leftarrow (val, dfltNtf)\}$ /* access to p_i's config */
22 **macro** $increment(phs) = \{$**select**(phs) **case** 0**: return** 0; **case** 1**: return** 2; **case** 2**: return** 0;$\}$
23 **macro** $allSeen() = (\text{FD}[i].part \subseteq (\text{allSeen} \cup \{p_i : \text{all}[i]\}))$;
24 **do forever begin**
25 **foreach** $p_k \notin \text{FD}[i].part$ **do** $(\text{config}[k], \text{prp}[k]) \leftarrow (\sharp, dfltNtf)$; /* clean after crashes */ **if** $((\exists p_k{:}((\text{prp}[k] = \langle 0, s\rangle) \wedge (s \neq \perp)) \vee (\text{config}[k] \in \{\perp, \emptyset\}))) \vee (\nexists p_k, p_{k'} \in \text{FD}[i].part{:}|degree(k) - degree(k')| > 1) \vee (\exists p_k \in \text{FD}[i].part{:}((\text{prp}[i].phase = x) \wedge (\text{prp}[k].phase = (x{+}1)(\bmod 3)) \wedge (p_k \notin allSeen) \wedge (x \in \{1,2\})) \vee (|notifSet| > 1) \wedge ((\{(\text{FD}[i], \text{FD}[i].part)\} = \{(\text{FD}[k], \text{FD}[k].part)\}_{p_k \in \text{FD}[i].part}) \wedge ((\text{config}[i] \cap \text{FD}[i].part) = \emptyset)))$ **where** $notifSet = \{\text{prp}[k].set{:}\exists p_{k'} \in \text{FD}[i].part{:}\text{prp}[k'] = \langle 2, \bullet\rangle\}_{p_k \in \text{FD}[i].part}$ **then** $configSet(\perp)$ **if** $(maxNtf() = \perp)$ **then** /* when no notification arrived */
26 **if** $|\{\text{config}[k]\}_{p_k \in \text{FD}[i]} \setminus \{\perp, \sharp\}| > 1$ **then** $configSet(\perp)$; /* nullify the configuration upon conflict */ **if** $(\text{config}[i] = \perp \wedge |\{\text{FD}[j] : p_j \in \text{FD}[i]\}| = 1)$ **then** $configSet(\text{FD}[i])$; /* reset during admissible runs */
27 **else**
28 all[i] $\leftarrow \bigwedge_{p_k \in \text{FD}[i].part}(echoNoAll(k) \wedge same(k))$; /* test all-the-same reports */ **foreach** $p_k \in \text{FD}[i].part : (echoNoAll(k) \wedge same(k))$ **do** allSeen $\leftarrow$ allSeen $\cup \{p_k\}$ **if** $echo() \wedge allSeen()$ **then** $(\text{prp}[i].phase, \text{allSeen}) \leftarrow (increment(\text{prp}[i].phase), \emptyset)$ {**select**(prp[i]$.phase$) **case** 0**:** prp[i] $\leftarrow dfltNtf$, **case** 1**:** prp[i] $\leftarrow maxNtf()$, **case** 2**:** config[i] $\leftarrow$ prp[i]$.set$};
29 **if** config[i] $\neq \sharp$ **then foreach** $p_j \in \text{FD}[i]$ **do send**$(\langle \text{FD}[i], \text{config}[i], \text{prp}[i], \text{all}[i], (\text{FD}[j].part, \text{prp}[j], \text{all}[j])\rangle)$
30 **upon receive** $m = \langle \text{FD}, \text{config}, \text{prp}, \text{all}, \text{echo}\rangle$ **from** p_j **do** $(\text{FD}[j], \text{config}[j], \text{prp}[j], \text{all}[j], \text{echo}[j]) \leftarrow m$;
31 **upon interrupt p_i's booting do foreach** p_k **do** $\text{echo}[k] \leftarrow (\text{config}[k], \text{prp}[k], \text{all}[k]) \leftarrow (\sharp, dfltNtf, false)$

monitor the intersection between the current configuration and the set of active participants (line 25). In case it is empty, the processors (participants or not) essentially begin a brute-force stabilization (outlined

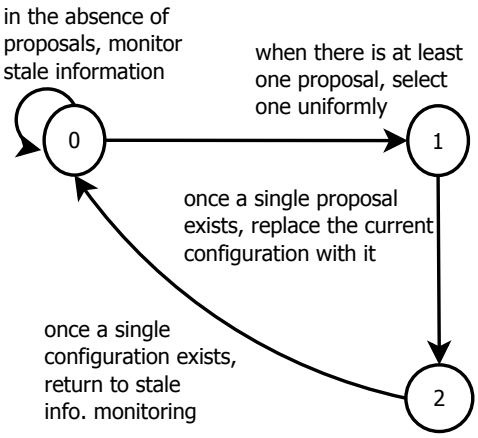

Figure 2: The configuration replacement automaton

below) where there is no more blocking of joining processors to become participants.

Brute-force stabilization. The algorithm detects the presence of stale information and recovers from these transient faults. *Configuration conflicts* are one of several kinds of such stale information and they refer to differences in the field config, which stores the (quorum) configuration values. Processor p_i can signal to all processors that it had detected stale information by assigning $\perp$ to config$_i$ and by that start a reset process that nullifies all config fields in the system (lines 25 and 26). Algorithm 3.1 uses the brute-force technique for letting processor p_i to assign to config$_i$ its set of trusted processors (line 26), which the failure detector FD$_i$ at processor p_i provides. Note that by the end of the brute-force process, all active processors (joining or participant) become participants. We show that eventually all active processors share identical (non-$\perp$) config values by the end of this process.

Delicate (configuration) replacement. Participants can propose to replace the current configuration with a new one, set, via the $estab$(set) interface. This replacement uses the *configuration replacement* process, which for the purposes of the overview, we abstract it as the automaton depicted in Figure 2. When the system is free from stale information, the configuration uniformity invariant (of the config field values) holds. Then, any number of calls to the $estab$(set) interface starts the configuration replacement process, which controls the configuration replacement using the following three phases: (1) selecting (deterministically and uniformly) a single proposal (while verifying the eventual absence of "unselected" proposals), (2) replacing (deterministically and uniformly) all config fields with the jointly selected proposal, and (3) bringing back the system to a state in which it merely tests for stale information.

The configuration replacement process requires coordinated phase transition. Algorithm 3.1 lets processor p_i to represent proposals as prp$_i[j] = (phase, set)$, where p_j is the processor from which p_i received the proposal, $phase \in \{0, 1, 2\}$ and set is a processor set or the null value, $\perp$. The *default proposal*, $\langle 0, \perp \rangle$, refers to the case in which prp encodes "no proposal". When p_i calls the function $estab$(set), it changes prp to $\langle 1, set \rangle$ (line 13) as long as p_i is not aware of an ongoing configuration replacement process, i.e., $noReco()$ returns true. Upon this change, the algorithm disseminates prp$_i[i]$ and by that guarantees eventually that $noReco()$ returns false for any processor that calls it. Once that happens, no call to $estab$(set) adds a new proposal for configuration replacement and no call to $participate()$ lets a joining processor to become a participant (line 14). The algorithm can then use the lexical value of the prp$_i[]$'s tuples for selecting one of them deterministically (Figure 2). To that end, each participant makes sure that all other participants report the same tuples by waiting until they "echo" back the same values as the ones it had sent to them. Once that

happens, the participant p_i makes sure that the communication channels do not include other "unselected" proposals by raising a flag ($\text{all}_i = true$) and waiting for the echoed values of these three fields, i.e., participant set, $\text{prp}_i[i]$ and all_i. This waiting continues until the echoed values match the values of any other active participant in the system (while monitoring their well-being). Before this participant proceeds, it makes sure that all active participants have noticed its phase completion (line 28). Each processor p maintains the allSeen variable; a set of participants that have noticed p's phase completion (line 28) and are thus added to p's allSeen set.

The above mechanism for phase transition coordination allows progression in a unison fashion. Namely, no processor starts a new phase before it has seen that all other active participants have completed the current phase and have noticed that all other have done so (because they have identical participants' set, prp and all[] values). This is the basis for emulating every step of the configuration replacement process (line 28) and making sure that the phase 2 replacement occurs correctly before returning to phase 0, in which the system simply tests for stale information. We show that since the failure detectors monitor the participants' well-being, this process terminates.

Detailed Descripion

We now proceed to a detailed, line-by-line description of Algorithm 3.1.

Variables. The algorithm uses a number of fields that each active participant broadcasts to all other system processors. The processor stores the values that they receive in arrays. Namely, we consider both p_i's own value (the i-th entry) and p_j's most recently received value (the j-th entry).

- The field config[]: an array in which config$[i]$ is p_i's view on the current configuration. Note that p_i assigns the *empty (configuration)* value $\perp$ after receiving a conflicting (not the same) non-empty configuration value, i.e., the received configuration is different than p_i's configuration. We use the symbol $\sharp$ for denoting that processor p_i is not a participant, i.e., $\text{config}_i[i] = \sharp$.

- The field FD[] is an array in which FD$[i]$ represents p_i's failure detector of trusted processors (without the list of processors that are suspected to be crashed).

- FD[]$.part$ is the participants' set, where FD$[i].part$ is an alias for $\{p_j \in \text{FD}[i] : \text{config}[j] \neq \sharp\}$ and FD$[j].part$ refers to the last value received from p_j. Namely, the FD field of every message encodes also this participation information.

- The field prp[] is an array of pairs $\langle phase \in \{0,1,2\}, set \subseteq P\rangle$, where prp$[i]$ refers to p_i's configuration replacement notifications. In the pair prp$[i]$, the field set can either be $\perp$ ('no value') or the proposed processor set.

- The field echo[] is an array in which echo$[i]$ is an alias of (FD$[i].part$, prp$[i]$, all$[i]$) and echo$[j]$ refers to the most recent value that p_i has received from p_j after p_j had responded to p_i with the most recent values it got from p_i.

- The field all[] is an array of Booleans, where all$[i]$ refers to the case in which p_i observes that all trusted processors have noticed its current (maximal) notification and they hold the same notification.

- The variable allSeen is a set that includes the processors p_k for which p_i received the all$[k]$ indication.

Constants, functions and macros. The constant $dfltNtf$ (line 15) denotes the default notification tuple $\langle 0, \perp \rangle$. The following functions define the interface between Algorithms 3.1 and 3.2 (Reconfiguration Management layer) and the Joining Mechanism (Algorithm 3.3). Note that the behavior which we specify below considers legal executions.

- The function $chsConfig()$ (line 10) returns config whenever there is a single such non-$\sharp$ value. Otherwise, $\perp$ is returned.
- The function $getConfig()$ (line 11) allows Algorithms 3.2 and 3.3 to retrieve the value of the current quorum configuration, i.e., config$_i[i]$. We note that during legal executions, this value is a set of processors whenever p_i is a participant. However, this value can be $\sharp$ whenever p_i is not a participant and $\perp$ during the process of configuration reset.
- The function $noReco()$ (line 12) returns $true$ whenever (1) p_i was not recognized as a trusted processor by a processor that p_i trusts, (2) there are configuration conflicts, (3) the participant sets have yet to stabilize, (4) there is an on-going process of brute force stabilization, or (5) there is a delicate (configuration) replacement in progress. This part of the interface allows Algorithms 3.2 and 3.3 to test for the presence of local evidence according to which Algorithm 3.1 shall disable delicate (configuration) replacement and joining to the participant set.
- The function $estab$(set) (line 13) provides an interface that allows the *recMA* layer to request from Algorithm 3.1 to replace the current quorum configuration with the proposed set, which is a non-empty group of participants. Note that Algorithm 3.1 disables this functionality whenever $noReco() = false$ or $set =$ config$[i]$.
- The function $participate()$ (line 14) provides an interface that allows the Joining Mechanism to request from Algorithm 3.1 to let p_i join the participant set, which is the group that can participate in the configuration and request the replacement of the current configuration with another (via the $estab$(set) function). Note that Algorithm 3.1 disables this functionality whenever $noReco() = false$ and thus there exists a single configuration in the system, i.e., the call to $chsConfig()$ (line 10) returns the single configuration that all active participants store as their current quorum configuration. This is except for case in which $(\{\text{config}_i[k]\}_{p_k \in \text{FD}_i[i]} \setminus \{\sharp\}) = \emptyset$. Here, $chsConfig()$ returns $\perp$, which starts a reset process in order to deal with a complete collapse where the quorum system includes no active participants.

Algorithm 3.1 uses the following macros.

- The macro $degree(k)$ (line 16) calculates the degree of p_k's most-recently-received notification degree, which is twice the notification phase plus one whenever all participants are using the same notification (and zero otherwise).
- The macros $echoNoAll(k)$ and $echo(k)$ (lines 17, and 18 respectively) test whether p_i was acknowledged by all participants for the values it has sent. The former function considers just the fields that are related to its own participants' set and notification, whereas the latter considers also the field all[].
- The macro $same(k)$ (line 19) performs a similar tests to the one of $echoNoAll(k)$, but considers only processor p_k's most-recently-received values rather than all participants' echoed ones.

- The macro $maxNtf()$ (line 20) selects a notification with the maximal lexicographical value or returns $\perp$ in the absence of notification that is not the default (phase 0) notification. We define our lexicographical order between prp_1 and prp_2, as $\text{prp}_1 \leq_{lex} \text{prp}_2 \iff ((\text{prp}_1.phase < \text{prp}_2.phase) \vee ((\text{prp}_1.phase = \text{prp}_2.phase) \wedge (\text{prp}_1.set \leq_{lex} \text{prp}_2.set)))$, where $\text{prp}_1.set \leq_{lex} \text{prp}_2.set$ can be defined using a common lexical ordering and by considering sets of processors as ordered tuples that list processors in, say, an ascending order with respect to their identifiers.
- The macro $configSet(val)$ (line 21) acts as a wrapper function for accessing p_i's local copies of the field config. This macro also makes sure that there are no (local) active notifications.
- The macro $increment(phs)$ (line 22) performs the transition between the phases of the delicate configuration replacement technique.
- The macro $allSeen()$ (line 23) tests whether all active participants have noticed that all other participants have finished the current phase.

The do forever loop. A line-by-line walkthrough of the pseudocode of Algorithm 3.1 follows.

Cleaning up, removal of stale information and invariant testing. The do forever loop starts by making sure that non-participant nodes cannot have an impact on p_i's computation (line 25) before testing that p_i's state does not include any stale information (line 25). Algorithm 3.1 tests for several types of stale information (c.f. Definition 3.1): (type-1) notifications in phase 0 must not have a $set \neq \emptyset$, (type-2) configurations that refer to the empty set, execute configuration reset, or conflicting configurations, (type-3) the degree gap between two notifications is greater than one, there are participants in different phases but the one in the more advanced phase does not appear in the $notifSet$ set, or the local set of notifications includes more than one notification and at least one of them is in phase 2, and (type-4) the quorum configuration includes at least one active participant (to avoid false positive error by testing only when processor p_i has a stable view of set of trusted processors and participants). In case any of these tests fails, the algorithm starts a configuration reset process by calling $configSet(\perp)$.

The brute-force stabilization technique. As long as no active notifications are present locally (line 25), every processor performs this technique for transient fault recovery. In the presence of configuration conflicts, the algorithm starts the configuration reset process (line 26). Moreover, during the configuration reset process, the algorithm waits until all locally trusted processors report that they trust the same set of processors (line 26).

The delicate replacement technique — phase synchronization. This technique synchronizes the system phase transitions by making sure that all active participants work with the same notification.

Each active participant tests whether all other trusted participants have echoed their current participants' set and notifications and have the same values with respect to these two fields (line 28). The success of this test assigns $true$ to the field $\text{all}_i[i]$. The algorithm then extends this test to include also the field all[] (line 28). Upon the success of this test with respect to participant p_k, the algorithm adds p_k to the set allSeen_i. Once processor p_i receives reports from all participants that the current phase is over, it moves to the next phase (line 28).

The delicate replacement technique — finite-state-machine emulation. Each of the three phases represent an automaton state (line 28); recall Figure 2 and the configuration replacement process discussed in the Overview. During phase 1, the system converges to a single notification. During phase 2, the system replaces the current configuration with the proposed one. Next, the system returns to its ideal state, i.e., phase 0, which allows new participants to join, as well as further reconfigurations.

Message exchange and the control of newly arrived processors. When a participant finishes executing the do forever loop, it broadcasts its entire state (line 29). Once these messages arrive, processor p_i stores them (line 30). The only way for a newly arrived processor to start executing Algorithm 3.1 is by responding to an interrupt call (line 31). This procedure notifies the processor state and makes sure that it cannot broadcast messages (line 29). The safe entry of this newly arrived processor to the participants' set, is via the function $participate()$ (line 14), which enables p_i's broadcasting. Thus, non-participants merely follow the system messages until the function $noReco()$ returns true and allows its join to the participant set by a call (from the Joining Mechanism) to the function $participate()$ (line 14).

3.1.2 Correctness

We first provide an outline of the proof and then proceed in steps to establish the correctness of the algorithm.

Outline

We say that system state c has no stale information when (1) all the notifications are valid, (2) there are no configuration conflicts or active reset process, (3) the phase information (including the set $allSeen$) are not out of synch, and (4) there are active participants in config. The correctness proof of Algorithm 3.1 shows that eventually there is no stale information, because they are all cleaned (line 25) or detected and cause configuration reset by calling $configSet(\perp)$ (line 25).

In the absence of notifications (line 25), Algorithm 3.1 merely tests for configuration conflicts (line 26) and during a configuration reset, it waits until all locally trusted processors report that they trust the same processors (line 26) before these processors become the configuration. I.e., it applies the brute force stabilization technique. The proof here shows that during admissible executions, the $\perp$ symbol propagates to all of the config fields in the system until all active processors assign to config either $\perp$ or the entire set of trusted processors. This reset process ends when all processors assign merely the latter value to config.

In the presence of notifications, Algorithm 3.1 synchronizes the system phase transitions by making sure that all active participants work with the same notification (lines 28 to 28) before moving to the next phase (line 28). The algorithm's three phases represent an automaton state (line 28). During phase 1, the system converges to a single notification. While $phase = 2$ holds, the system replaces the current configuration with the proposed one. Next, the system returns to its no-notifications state, i.e., $phase = 0$, which allows new participants as well as further reconfigurations. When a participant finishes the do forever loop, it broadcasts its entire state (line 29). Once these messages arrive, the receiving processor stores them (line 30).

The only way for a newly arrived processor to start executing Algorithm 3.1 is by responding to an interrupt call (line 31). This procedure nullifies the state of the newly arrived processor with config $= \sharp$ and by that makes sure that it cannot broadcast (line 29). The safe entry of this processor to the participating set, is via the function $participate()$ (line 14), which enables p_i's broadcasting. Thus, non-participants merely follow the system messages until the function $noReco()$ returns true and allows its join to the participant set by a call to $participate()$ (line 14). By controlling the new joins, the three phase structure is the basis

for proving the final theorem stating that eventually a common configuration set is adopted by every active processor. We now proceed to the detailed proof.

Configuration conflicts and stale information

We begin by classifying the stale information into four types.

Definition 3.1 *We say that processor p_i in system state c has a stale information in c of:*

type-1: *when* $(\exists p_k : ((\text{prp}[k] = \langle 0, s\rangle) \wedge (s \notin \{\emptyset, \perp\})))$ *(cf. line 25).*

type-2: *when* $(\exists p_k : (\text{config}[k] \in \{\perp, \emptyset\}))$ *(cf. line 25) or c encodes a* (configuration) conflict*, i.e., there are two active processors p_i and p_j for which* $\text{config}_i[i] \neq \text{config}_j[j]$, $\text{config}_i[i] \neq \text{config}_i[j]$*, or* $m_{j,i}.\text{config} \neq \text{config}_i[i]$ *in any message in the communication channel from p_i to p_j.*

type-3: *when* $(\nexists p_k, p_{k'} \in \text{FD}_i[i].part : |degree_i(k) - degree_i(k')| > 1) \vee (\exists p_k \in \text{FD}_i[i] : ((\text{prp}_i[i].phase = x) \wedge (\text{prp}_i[k].phase = (x+1) (\text{mod } 3)) \wedge (p_k \notin allSeen_i) \wedge (x \in \{1, 2\})) \vee (|notifSet_i| > 1))$*, where* $notifSet_i = \{\text{prp}_i[k].set : \exists p_{k'} \in \text{FD}_i[i].part : \text{prp}_i[k'] = \langle 2, \bullet\rangle\}_{p_k \in \text{FD}_i[i]}$.

type-4: *when* $((\{(\text{FD}_i[i], \text{FD}_i[i].part)\} = \{(\text{FD}_i[k], \text{FD}_i[k].part)\}_{p_k \in \text{FD}_i[i].part}) \wedge ((\text{config}_i[i] \cap \text{FD}_i[i].part) = \emptyset))$.

Claim 3.1 (Eventually there is no type-1 stale information) *Eventually the system reaches a state $c \in R$ in which the invariant of no type-1 stale information holds thereafter.*

Proof. Let $c \in R$ be a system state in which processor p_i has an applicable step a_i that includes the execution of the do forever loop (line 24 to 29). We note that immediately after a_i, processor p_i has no type-1 stale information (line 25 removes it). Moreover, that removal always occurs before a_i sends any message m (line 29). Therefore, eventually for every active processor p_j that receives from p_i message m (line 30), it holds that $(m.phase = 0) \iff (m.\text{prp} = (0, \perp))$ as well as for every item $\text{prp}_i[k] : p_k \in P$. Once this invariant holds for every pair of active processors p_i and p_j, the system reaches state c. We conclude the proof by noting that Algorithm 3.1 never assigns to $\text{prp}_i[j]$ a values that violates this claim invariant. □

Replacement state and message; explicit and spontaneous replacements

We say that a processor p_i's state encodes a *(delicate) replacement* when $\text{prp}_i[j] \neq \langle 0, \perp\rangle$ and we say that a message $m_{i,j}$ in the channel from p_i to p_j encodes a *(delicate) replacement* when its $\text{prp} \neq \langle 0, \perp\rangle$. Given a system execution R, we say that R *does not include an explicit (delicate) replacement* when throughout R no node p_i calls $estab()$. Suppose that execution R does not include an explicit (delicate) replacement and yet there is a system state $c \in R$ in which a processor state or a message in the communication channels encodes a (delicate) replacement. In this case, we say that R includes a *spontaneous (delicate) replacement*.

Lemma 3.2 (Eventually there is no type-2 stale information) *Let R be an admissible execution that does not include explicit (delicate) replacements nor spontaneous ones. The system eventually reaches a state $c \in R$ in which the invariant of no type-2 stale information holds thereafter.*

Proof. Note that any of R's steps that includes the do forever loop (line 24 to 29) does not run lines 28 to 28 (since R does not include an explicit nor spontaneous replacement). If R's starting system state does not include any configuration conflicts, we are done. Suppose that R's starting system state does include a conflict, i.e., $\exists p_i, p_j \in P : (\text{config}_i[i] = \perp) \vee (\text{config}_i[i] \neq \text{config}_i[j]) \vee (\text{config}_i[i] \neq \text{config}_j[j])$ or there

is a message, $m_{i,j}$, in the communication channel from p_i to p_j, such that the field $(m_{i,j}.\text{config}[k] = \perp) : p_k \in \text{FD}_i[i] \vee (m_{i,j}.\text{config} \neq \text{config}_i[i])$, where both p_i and p_j are active processors. In Claims 3.3, 3.4 and 3.5 we show that in all of these cases, eventually $\forall p_i \in P : \text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ before showing that eventually there are no configuration conflicts (Claim 3.6).

Claims 3.3, 3.4 and 3.5 consider the values in the field config that are either held by an active processor $p_i \in P$ or in its outgoing communication channel to another active processor $p_j \in P$. We define the set $S = \{S_i \cup S_out_i\}_{p_i \in P}$ to be the set of all these values, where $S_i = \{\text{config}_i[j]\}_{p_j \in \text{FD}_i[i]}$ and $S_out_i = \{m_{i,j}.\text{config}\}_{p_j \in \text{FD}_i[i]}$.

Claim 3.3 *Suppose that in R's starting system state, there are processors $p_i, p_j \in P$ that are active in R, for which $|S \setminus \{\perp, \sharp\}| > 1$, where $\exists S' \subseteq S : S' \in \{\{\text{config}_i[i], \text{config}_i[j]\}, \{\text{config}_i[i], m_{i,j}.\text{config}\}\}$. Eventually the system reaches a state in which $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ holds.*

Proof. Suppose that $S' = \{\text{config}_i[i], \text{config}_i[j]\}$ holds. Immediately after R's starting state, processor p_i has an applicable step that includes the execution of the do forever loop (line 24 to 29). In that step, the if-statement condition $(|\{\text{config}_i[k] : p_k \in \text{FD}_i[i]\} \setminus \{\perp, \sharp\}| > 1)$ (line 26's if-statement) holds, p_i assigns $\perp$ to $\text{config}_i[i]$ and the proof is done. Suppose that $S' = \{\text{config}_i[i], m_{i,j}.\text{config}\}$ holds. Upon $m_{i,j}$'s arrival, processor p_i assigns $m_{i,j}.\text{config}$ to $\text{config}_i[j]$ (line 30) and the case of $S' = \{\text{config}_i[i], \text{config}_i[j]\}$ holds. □

Claim 3.4 *Suppose that in R's starting system state, there are processors $p_i, p_j \in P$ that are active in R, for which $|S \setminus \{\perp, \sharp\}| > 1$, where $\exists S' \subseteq S : S' \in \{\{\text{config}_i[i], \text{config}_j[j]\}\}$. Eventually the system reaches a state in which $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ or $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$ holds.*

Proof. Suppose, towards a contradiction, for any system state $c \in R$ that neither $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ nor $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$. Note that p_i and p_j exchange messages eventually, because whenever processor p_i repeatedly sends the same message to processor p_j, it holds that p_j receives that message eventually (the fair communication assumption, Section 2) and vice versa. Such message exchange implies that the case of $|S \setminus \{\perp, \sharp\}| > 1$ (Claim 3.3) holds eventually, where $\exists S' \subseteq S : S' \in \{\{\text{config}_i[i], m_{i,j}.\text{config}\}, \{\text{config}_j[j], m_{i,j}.\text{config}\}\}$. Thus, we reach a contradiction and therefore eventually the system reaches a state in which $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ or $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$ hold. □

Claim 3.5 *Suppose that in R's starting system state, there is a processor $p_i \in P$ that is active in R, for which $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$. (1) For any system state $c \in R$, it holds that $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$. Moreover, (2) $R = R' \circ R''$ has a suffix, R'', for which $\forall c'' \in R'' : \forall p_i, p_j : (\{m_{i,j}.\text{config}, \text{config}_j[i], \text{config}_j[j]\} \setminus \{\perp, \text{FD}_i[i]\}) = \emptyset$.*

Proof. We prove each part of the statement separately.
Part (1). We start the proof by noting that $\forall p_i, p_j \in P$, it holds that, throughout R, we have that $\text{FD}_i[i]$'s value does not change and that $\text{FD}_i[i] = \text{FD}_j[j]$, because this lemma assumes that R is admissible. To show that $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$ holds in any $c \in R$, we argue that any step $a_i \in R$ in which p_i changes $\text{config}_i[i]$'s value includes the execution of line 26 or line 26 (see the remark at the beginning of this lemma's proof about $a_i \in R$ not including the execution of lines 28 to 28), which assign to $\text{config}_i[i]$ the values $\perp$, and respectively, $\text{FD}_i[i]$.
Part (2). In this part of the proof, we first consider the values in $m_{i,j}.\text{config}$ and $\text{config}_j[i]$ before considering the one in $\text{config}_j[j]$.
Part (2.1). To show that in $c'' \in R''$ it holds that $\forall p_i, p_j : \{m_{i,j}.\text{config}, \text{config}_j[i]\} \setminus \{\perp, \text{FD}_i[i]\} = \emptyset$, we note that when p_i loads a message $m_{i,j}$ (line 29) before sending to processor p_j, it uses $\text{config}_i[i]$'s value for

the field config. Thus, eventually $m_{i,j}$.config $\in \{\perp, \text{FD}_i[i]\}$ and therefore $\text{config}_j[i] \in \{\perp, \text{FD}_i[i]\}$ records correctly in $c'' \in R''$ the most recent $m_{i,j}$'s value that p_j receives from p_i (line 30).

Part (2.2). To show that eventually $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$, we note that once p_j changes the value of $\text{config}_j[j]$, it holds that $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$ thereafter (due to the remark in the beginning of this lemma, which implies that only lines 26 and 26 can change $\text{config}_j[j]$, and by the part (1) of this claim's proof while replacing the index i with j). Suppose, towards a contradiction, that p_j does not change that value of $\text{config}_j[j]$ throughout R and yet $\text{config}_j[j] \notin \{\perp, \text{FD}_i[i]\}$. Note that $(|\{\text{FD}[j] : p_j \in \text{FD}[i]\}| = 1)$ (see the second clause of the if-statement condition in line 26) holds throughout R, because R is admissible. Therefore, whenever p_i takes a step that includes the execution of the do forever loop (line 24 to 29), processor p_i sends to p_j the message $m_{i,j}$, such that $m_{i,j}$.config $= \text{config}_i[i]$ (line 29) and $\text{config}_i[i] = \text{FD}_i[i]$ (see part (2.1) of this proof). Since p_i sends $m_{i,j}$ repeatedly, processor p_j receives eventually $m_{i,j}$ (the fair communication assumption, Section 2) and stores in $\text{config}_j[i] = m_{i,j}.\text{config} = \text{config}_i[i] = \text{FD}_i[i] \neq \perp$. Immediately after that step, the system state allows p_j to take a step in which the condition $(|\{\text{config}_j[k] : p'_j \in \text{FD}_j[k]\} \setminus \{\perp, \sharp\}| > 1)$ (line 26's if-statement) holds and p_j changes $\text{config}_j[j]$'s value to $\perp$. Thus, this part of the proof end with a contradiction, which implies that the system reaches a state in which $\text{config}_j[j] \in \{\perp, \text{FD}_i[i]\}$. □

Claim 3.6 *Suppose that in R's starting system state, it holds that for every two processors $p_i, p_j \in P$ that are active in R, we have that $(\{\text{config}_i[i], \text{config}_j[i], m_{i,j}.\text{config}\} \setminus \{\perp, \text{FD}_i[i]\}) = \emptyset$, where $m_{i,j}$ is a message in the channel from p_i to p_j. Eventually the system reaches a state in which* $\text{config}_i[i] = \text{FD}_i[i]$.

Proof. By this claim assumptions, we have that in R's starting system state, the if-statement condition $(|\{\text{config}_i[k] : p_k \in \text{FD}_i[k]\} \setminus \{\perp, \sharp\}| > 1)$ (line 26) does not hold. Moreover, $|\{\text{FD}_i[j] : p_j \in \text{FD}_i[i]\}| = 1$ (line 26) holds throughout R, because R is admissible. Therefore, this claim assumptions with respect to R's starting states actually hold for any system state $c \in R$, because only lines 26 and 26 can change the value of $\text{config}_i[i] \in \{\perp, \text{FD}_i[i]\}$, which later p_i uses for sending the message $m_{i,j}$ (line 29), and thus also $m_{i,j}$.config $\in \{\perp, \text{FD}_i[i]\}$ as well as $\text{config}_j[i] \in \{\perp, \text{FD}_i[i]\}$ records correctly the most recent $m_{i,j}$'s that p_j receives from p_i (line 30).

To conclude this proof, we note that immediately after any system state $c \in R$, processor p_i has an applicable step $a_i \in R$ that includes the execution of line 26 (by similar arguments to the ones used by the first part of this proof). Moreover, a_i does not include the execution of line 26, because by the first part of this proof the condition of the if-statement of line 26 does not hold. In the system state that immediately follow a_i, the invariant $\text{config}_i[i] = \text{FD}_i[i]$ holds. □

By this lemma's assumption, there is no configuration $c \in R$ replacement state nor replacement message. Claim 3.6 implies that the system reaches a state $c_{noConf} \in R$ that has no configuration conflict eventually. Thus, c_{noConf} is safe. □

Claim 3.7 (Eventually there is no type-4 stale information) *Let R be an admissible execution of Algorithm 3.1. Eventually the system reaches a state $c \in R$ in which the invariant of no type-4 stale information holds thereafter.*

Proof. Without the loss of generality, suppose that there is no system state in R that encodes a configuration conflict. (We can make this assumption without losing generality because Lemma 3.2 implies that this claim is true whenever this assumption is false.) Moreover, let $c \in R$ be a system state in which processor p_i has an applicable step a_i that includes the execution of the do forever loop (line 24 to 29).

Since R is admissible, $(\{(\text{FD}_i[i], \text{FD}_i[i].part)\} = \{(\text{FD}_i[k], \text{FD}_i[k].part)\}_{p_k \in \text{FD}_i[i].part})$ holds in c. Therefore, the case in which $((\text{config}_i[i] \cap \text{FD}_i[i].part) = \emptyset)$ in c implies a call to the function $configSet(\perp)$ (line 25) in the step that immediately follows. By using Lemma 3.2, we have that this lemma is true. □

Phase and degree progressions

Let R be an execution of Algorithm 3.1 that is admissible with respect to the participant sets. Suppose that p_i is a processor that is active in R and that $p_i \in \text{FD}[i].part$. We say that processor p_i is an active participant in R. Let $p_i, p_j, p_k \in P$ be processors that are active participants in R and $c \in R$ a system state. We denote by $\text{NA}(c) = \{(\text{prp}_j[k], \text{all}_j[k])\}_{p_j,p_k \in \text{FD}_i[i]} \cup \{(\text{prp}_j[k], \text{all}_j[k]) : m_{j,k} = \langle \bullet, \text{prp}_j[k], \text{all}_j[k], (\bullet)\rangle \in channel_{k,j}\}_{p_j,p_k \in \text{FD}_i[i]} \cup \{(\text{echo}_j[k].\text{prp}, \text{echo}_j[k].\text{all}) : m_{j,k} = \langle \bullet, (\bullet, \text{echo}_j[k].\text{prp}, \text{echo}_j[k].\text{all})\rangle \in channel_{k,j}\}_{p_j,p_k \in \text{FD}_i[i]} \setminus \{(\langle 0, \perp\rangle, \bullet)\}$ the set of all pairs that includes the notification and all fields that appear in c (after excluding the default notification, $\langle 0, \perp\rangle$, while including all the information in the processors' states and communication channels as well as their replications, e.g., the echo field). We denote by $N(c) = \{\text{n} : (\text{n}, \text{a}) \in \text{NA}(c)\}$ the set of notifications that appear in c. We denote the *degree set* of notification $\text{n} \in N(c)$ by $D(c, \text{n}) = \{2 \cdot \text{n}.phase + |\{1 : \text{a}\}| : (\text{n}, \text{a}) \in \text{NA}(c)\}$. For a given system state $c \in R$, notification $\text{n} \in N(c)$ that appears in c, we denote by $S(c, \text{n}) = \{\text{n}' : ((\text{n}' \in N(c)) \wedge (\text{n}.set = \text{n}'.set))\}$ the set of all notifications $\text{n}' \in N(c)$ that have the same set filed as the one of n and $S(c) = \{\text{n}.set : \text{n} \in N(c)\}$ is the set of all notification sets in c.

Lemma 3.8 *Let R be an execution of Algorithm 3.1 that is admissible with respect to the participant sets and that does not include an explicit (delicate) replacement. Suppose that in R's starting system state, c, there are notifications, i.e., $N(c) \neq \emptyset$. Let $\text{n} \in N(c)$ be a notification for which it holds that $\forall \text{n}' \in N(c) : \text{n}' \leq_{\text{lex}} \text{n}$ in c (recall the definition of $\leq_{\text{lex}}$ in the description of the $maxNtf()$ macro). The system reaches eventually a state $c_{\not\exists \text{n}} \in R$ in which $S(c_{\not\exists \text{n}}, \text{n}) = \emptyset$.*

Proof. Let $a \in R$ be a step that immediately precedes the system state $c \in R$ and $c' \in R$ be the system state that immediately follows a. Note that when the system reaches a system state $c' \in R$ in which $\exists \text{n} : (S(c, \text{n}) \cap S(c', \text{n})) \neq \emptyset$, the proof is done. Suppose, in the way of a proof by contradiction, that this lemma is false, i.e., $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$, i.e., the notification n (and additional notifications that have the same proposed set but different phase number) does not "disappear" from the system after any step and "become" the default notification $\langle 0, \perp\rangle$. Claims 3.9, 3.10 and 3.12 are needed for the proof of Claim 3.13, which completes this proof by contradiction.

Claim 3.9 *Suppose that $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$. Without the loss of generality, we can assume that (1) no step in R includes a call to $configSet(\perp)$, (2) n's lexical value in R's starting system state, c, is the greater or equal to the lexical value of any notification in any system state $c'' \in R$, (3) $\text{prp}_i[i] = \text{n}$ in R's starting system state, where $p_i \in P$ is an active participant in R, and (4) step $a \in R$ that immediately follows R's starting system state includes the execution of the do forever loop (line 24 to 29).*

Proof. We prove each part of the claim separetly.
Part (1). This claim's assumption that $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$ implies that no step in R includes a call to $configSet(\perp)$, say, due to the if-statement condition of lines 25 or 26 hold in any system state $c'' \in R$. The reason is that Claim 3.4 implies that a call to $configSet(\perp)$ brings the system to a state $c''' \in R$ in which $S(c''', \text{n}) = \emptyset$. Namely, $c''' = c_{\not\exists \text{n}}$ and the proof of this lemma is done.
Part (2). We can choose a suffix of R, such that n's lexical value in R's starting system state, c, is the greater or equal to the lexical value of any notification in any system state $c'' \in R$. We can do that because there is a bounded number of possible lexical values and our assumption that $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$.
Part (3). We note that when considering the case in which c encodes a message $m_{i,j}$ that has n in one of its fields, message $m_{i,j}$ reaches eventually from p_j to p_i (the fair communication assumption, Section 2). Therefore, we suppose that notification n is encoded in p_i's state, i.e., $\text{prp}_i[j] = \text{n}$ in c, where $p_i, p_j \in P$ are

active participants in R. We can make this assumption, without the loss of generality, because we can take a suffix of R (that system reaches eventually) in which this assumption holds for its starting configuration and then we take that suffix to be the execution that this lemma considers, i.e., R. Using similar arguments about generality, we also assume that step a includes the execution of the do forever loop (line 24 to 29), and thus $\text{prp}_i[i] = \text{n}$ in c'.

Part (4). Use the same arguments as for part (3) of this proof. □

Claim 3.10 *Suppose that step a includes the execution of the do forever loop (line 24 to 29) immediately after and before the system states c, and respectively, c'. When $\text{prp}_i[i].set = \text{n}.set$ holds in c, we have that $\text{prp}_i[i].phase \geq \text{n}.phase$ does not decrease in c'.*

Proof. Since $\text{prp}_i[i].phase$ changes only in lines 28 and 28 (case 1), the assumption that $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$ in any $c'' \in R$ implies that $\text{prp}_i[i] = \text{n}$ in any $c'' \in R$ that is after c'. The reason is that line 28 indeed does not decrease $\text{prp}_i[i].phase$ (part (2) of Claim 3.9) and line 28 only decreases $\text{prp}_i[i].phase$ when assigning 0 (cf. case 0 of the function $increment()$). However, the latter cannot occur for any $p_i \in R$ that is an active participant in R and for which $\text{prp}_i[i] = \text{n}$ holds in c' (part (3) of Claim 3.9). □

Claim 3.11 *For every system state $c'' \in R$, it holds that $\text{prp}_i[i] = \text{n}$.*

Proof. Since $\text{prp}_i[i] = \text{n}$ holds in c (Claim 3.9), since $\text{prp}_i[i].phase$ is non-decreasing during step a (Claim 3.10), since n is lexically greater or equal than any other notification in any system state of R (Claim 3.9) and by this proof's assumption that $\nexists c_{\not\exists \text{n}} \in R : (S(c, \text{n}) \cap S(c_{\not\exists \text{n}}, \text{n})) = \emptyset$ holds, it is true that $\text{prp}_i[i] = \text{n}$ holds in every system state $c'' \in R$. □

Claim 3.12 *The following sequence of invariants is true.*

(1) *Suppose that $\text{prp}_i[i] = \text{n}$ holds in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_j \in P$ that is an active participant in R, it holds that $\text{prp}_j[i] = \text{n}$ and $\text{FD}_j[i] = \text{FD}_i$. Moreover, $\text{prp}_j[j] = \text{n}$ and $\text{FD}_j[j] = \text{FD}_i$ in c'' eventually.*

(2) *Suppose that invariant (1) holds in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_i \in P$ that is an active participant in R, it holds that $\text{echo}_i[j].\text{prp} = \text{n}$, $\text{echo}_i[j].\text{part} = \text{FD}_i[i].part$ and $\text{prp}_i[j] = \text{n}$ in c''.*

(3) *Suppose that invariants (1) and (2) hold in every system state $c \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_i \in P$ that is an active participant in R, it holds that $\text{all}_i[i] = true$ in c''.*

(4) *Suppose that invariants (1), (2) and (3) hold in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_j \in P$ that is an active participant in R, it holds that $\text{all}_j[i] = true$ in c''.*

(5) *Suppose that invariants (1) to (4) hold in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_j \in P$ that is an active participant in R, it holds that $\text{echo}_i[j] = (\text{FD}_i[i].part, \text{prp}_i[i], \text{all}_i[i])$ in c''.*

(6) *Suppose that invariants (1) to (5) hold in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for any $p_j \in P$ that is an active participant in R, it holds that $p_i \in allSeen_j$ in c''.*

(7) *Suppose that invariants (1) to (6) hold in every system state $c' \in R$. Eventually the system reaches a state $c'' \in R$, such that for all $p_i \in P$ that is an active participant in R, it holds that the if-statement condition of line 28 holds in c''.*

Proof. We prove items (1) to (7).
(1) Since p_i repeatedly sends message $m_{i,j}$ to every active processor $p_j \in \text{FD}_i[j]$ (line 29), where $m_{i,j} = \langle \bullet, \text{prp} = \text{n}, \bullet \rangle$, message $m_{i,j}$ arrives eventually to p_j (line 30 and the fair communication assumption, Section 2). This causes p_j to store n in $\text{prp}_j[i]$ (at least) as long as $\text{prp}_i[j] = \text{n}$, i.e., in every system state c'' that follows $m_{i,j}$ arrival to p_j. For the case of $\text{n}.phase = 1$, we have that $\text{prp}_j[j] = \text{n}$ (Claim 3.10 and line 28 (case 1) as well as the fact that n has the greatest lexicographic value in every $c' \in R$, cf. part (2) Claim 3.9). The same arguments as above imply that $\text{FD}_j[i] = \text{FD}_i$. Moreover, $\text{FD}_j[j] = \text{FD}_i$ is implied by the assumption that R is admissible with respect to the participant sets.

We show that $\text{prp}_j[j] = \text{n}$ for any $\text{n}.phase \in \{0, 1, 2\}$. By the assumptions made in the beginning of the proof of this lemma, we do not need to consider the case of $\text{n}.phase = 0$. Thus, we start by considering the case of $\text{n}.phase = 1$, we have that $\text{prp}_j[j] = \text{n}$ (Claim 3.10 and line 28 (case 1) as well as the fact that n has the greatest lexicographic value in every $c' \in R$, cf. part (2) Claim 3.9). We now turn to consider the case of $\text{n}.phase = 2$. I.e., while keeping in mind that $\text{prp}_j[i] = \text{n}$, we need to show that $\text{prp}_j[j] = \text{n}$. We do that by showing that whenever $\text{prp}_j[j] \neq \text{prp}_j[i]$, the if-statement condition of line 25 holds for any value of $\text{prp}_j[j].phase \in \{0, 1, 2\}$ and we focus on the case of $((\text{prp}_j[j].set \neq \text{prp}_j[i].set) \wedge (\text{prp}_j[j].phase = \text{prp}_j[i].phase = 2))$. Note that by showing that the if-statement condition of line 25 does hold in c'', we have reached a contradiction (part (1) of Claim 3.9's statement). For the case of $\text{prp}_j[j].phase \in \{0, 1\}$, the if-statement condition of line 25 holds in c'' because $\text{prp}_j[j] \neq \text{prp}_j[i]$ implies $\text{all}_j[j] = false$ (line 28) and the condition $|degree_j(j) - degree_j(i)| > 1$ (line 25 does hold due to the fact that $degree_j(j) \leq 2$ and $degree_j(i) \geq 4$). Moreover, for the case of $((\text{prp}_j[j].set \neq \text{prp}_j[i].set) \wedge (\text{prp}_j[j].phase = \text{prp}_j[i].phase = 2))$, the condition $(|notifSet_i| > 1)$ (line 25) does not hold, because $\{\text{prp}_j[j].set, \text{prp}_j[i].set\} \subseteq notifSet$ and $|\{\text{prp}_j[i].set, \text{prp}_j[j].set\}| = 2$, where $notifSet = \{\text{prp}_j[k].set : \exists p_k \in \text{FD}_j[j] : \text{prp}_j[i] = \langle 2, \bullet \rangle\}_{p_k \in \text{FD}_j[j]}$.
(2) By similar arguments to the ones that show the arrival of $m_{i,j}$ in part (1) of this proof, processor p_j repeatedly sends the message $m_{j,i} = \langle \bullet, \text{prp} = \text{n}, \bullet, \text{echo} = (\bullet, \text{n}, \bullet) \rangle$ to p_i, which indefinitely stores n in $\text{prp}_i[j]$ and $\text{echo}_i[j]$ (line 30). The same arguments as above imply that $\text{echo}_i[j].\text{part} = \text{FD}_i[i].part$.
(3) The system reaches a state in which the invariants of parts (1) and (2) hold for every $p_k \in \text{FD}_i[i]$. In that system state, it holds that $\text{prp}_i[i] = \text{n}$ (this claim assumption), as well as the fact that $\text{echo}_i[j].\text{prp} = \text{n}$ and $\text{prp}_i[j] = \text{n}$ (part (2) of this proof for showing that $\text{prp}_j[j] = \text{n}$ eventually and then using part (1) for showing that $\text{FD}_i[j] = \text{FD}_j[j]$ and $\text{prp}_i[j] = \text{n}$). In other words, $(\forall p_k \in \text{FD}[i].part : (echoNoAll(k) \wedge same(k))$ (cf. line 28). Let us look at the first step after that state in which processor p_i executes the do forever loop (lines 24 to 29). Immediately after that step, the system reaches a state in which $\text{all}_i[i] = true$ holds (line 28).
(4) By similar arguments that appear in parts (1) and (2) of this proof, processor p_i sends repeatedly the message $m_{i,j} = \langle \bullet, \text{prp} = \text{n}, \text{all} = true, \bullet) \rangle$ to processor p_j. Once message $m_{i,j}$ arrives eventually to p_j (line 30 and the fair communication assumption, Section 2), processor p_j stores $true$ in $\text{all}_j[i]$ (as well as the notification n in $\text{prp}_j[i]$). This holds for every system state c'' that follows $m_{i,j}$'s arrival to p_j.
(5) By similar arguments that appear in parts (1) and (2) of this proof, processor p_j sends repeatedly the message $m_{i,j} = \langle \bullet, \text{prp} = \text{n}, \text{all} = true, \text{echo} = (\text{FD}_i[i].part, \text{n}, true) \rangle$ to processor p_i. Once message $m_{j,i}$ arrives eventually to p_i (line 30 and the fair communication assumption, Section 2), processor p_i stores $(\text{n}, true)$ in $\text{echo}_i[j]$. This holds for every system state c'' that follows $m_{j,i}$'s arrival to p_j.
(6) We show that, once the invariants of parts (1) to (5) of this proof hold, the for-each condition of line 28 holds as well. Moreover, we show that assuming, towards a contradiction, that $p_i \notin \text{allSeen}_j$ implies that there is a step in R that includes a call to $configSet(\perp)$ (line 25), which is false according to part (1) of Claim 3.9's statement. By showing both the necessity and sufficiency conditions, we can conclude that

the assertion of this part in the claim even thought in Algorithm 3.1 processors p_i and p_j do not exchange information about allSeen$_j$ in order to validate that the set allSeen$_j$ includes p_i.

The for-each condition of line 28 condition requires that for any $p_j \in$ FD$_i[i].part$ to have $(echoNoAll(j) \wedge same(j))$, where the index k is substituted here with j. The following is true: prp$_i[j] =$ prp$_i[i] \wedge$ FD$_i[j].part =$ FD$_i[i].part \wedge$ echo$_i[j].part =$ FD$_i[i].part$ (part (1) of this proof and the assumption that R is admissible with respect to participant sets), echo$_i[j]$.prp $=$ prp$_i[i]$ (part (5) of this proof), and all$_i[j] = true$ (part (1) of this proof and then apply (4) for the case of prp$_j[i] =$ n).

We now turn to show that $p_j \notin$ allSeen$_i$ implies that there is a step in R that includes a call to $configSet(\perp)$ (line 25). By the assumptions made in the beginning of the proof of this lemma, we do not need to consider the case of n.$phase = 0$. Note that for n.$phase = x \in \{1, 2\}$ and $p_i \notin allSeen_j$, when the following condition holds $(\exists p_i \in$ FD$_j[j] : (($prp$_j[j].phase = x) \wedge (prp[k].phase = (x+1) (\mathrm{mod}\ 3)) \wedge (p_i \notin allSeen_j)))$, the condition $(p_i \notin allSeen_j)$ implies that the if-statement condition of line 25 holds in c'', and this contradicts part (1) of Claim 3.9's statement. Thus, $p_i \in allSeen_j$ in c''.

(7) Recall that part (5) of this proof says that echo$_i[j] = (\bullet, \text{n}, true)$. By taking prp$_j[j] =$ n from part (1) and applying it to the results of parts (1) to (6), we get that $p_j \in allSeen_i$ holds eventually, because this is true for every pair of processors $p_i, p_j \in P$ that are active participants in R. Let us consider a system state in which invariants (1) to (6) hold for any such pair of active participants p_i and p_j. This implies that $allSeen_i()$ holds (by the fact that part (4) eventually implies that $\forall p_k \in$ FD$_i[i] : p_k \in$ allSeen$_i$). Therefore, the if-statement condition of line 28 holds (by this lemma's assumption that R is admissible with respect to the participant sets and part (4) of this proof that shows $\forall p_k \in$ FD$_j[i] :$ all$_j[k] = true$). □

Claim 3.13 shows a contradiction to the assumption made in the beginning of the proof of Lemma 3.8. This contradiction completes the proof of this lemma.

Claim 3.13 *$\exists\, c_{\exists \text{n}} \in R$ in which $S(c_{\nexists \text{n}}, \text{n}) = \emptyset$.*

Proof. Part (7) of Claim 3.12 says that the if-statement condition of line 28 holds in $c'' \in R$ eventually. Let a_i a step that immediately follows c'' and in which p_i executes the do forever loop (line 24 to 29). Note that the case of n.$phase \in \{0\}$ contradicts this lemma's assumption that $c_{\nexists \text{n}} \in R$ in which $S(c_{\nexists \text{n}}, \text{n}) = \emptyset$. Suppose that n.$phase \in \{1\}$ in c''. The step a_i includes an execution of line 28 that increases the lexical value of n (contradicting part (2) of Claim 3.9). Suppose that n.$phase \in \{2\}$ in c''. The step a_i includes an execution of line 28 that changes n's value to $\langle 0, \perp \rangle$. Claim 3.12 implies that this holds for any $p_i \in P$ that is an active participant in R. This is a contradiction with the assumption made in the beginning of the proof of this lemma, i.e., $\exists\, c_{\nexists \text{n}} \in R$ in which $S(c_{\nexists \text{n}}, \text{n}) = \emptyset$. □

This completes the proof of Lemma 3.8. □

Algorithm 3.1 is self-stabilizing

Let $a_i \in R$ be a step in which processor p_i calls the function $estab$(set) (line 13), and in which the if-statement condition $(noReco() \wedge (set \notin \{$config$[i], \emptyset\}))$ does hold in the system state that immediately precedes a_i. We say that a_i is an *effective (configuration establishment)* step in R. Similarly, we consider $a_i \in R$ to be a step in which processor p_i calls the function $participate()$ (line 14), and in which the if-statement condition $noReco()$ does hold in the system state that immediately precedes a_i. Let $R = R' \circ R_{\text{VNER}} \circ R'''$ be an execution that does include explicit (delicate) replacements, where R' and R''' are a prefix, and respectively, a suffix of R. Let us consider R_{VNER}, which is a part of execution R. We say that R_{VNER} *virtually does not include explicit (delicate) replacements* (VNER) when for any step $a \in R_{\text{VNER}}$

that includes a call the function $estab$(set) (line 13) or $participate()$ (line 14) is ineffective. Given a system state $c \in R$, we say that c includes no notification if none of its active processors stores a notification and there are no notifications in transit between any two active processors.

Lemma 3.14 (Eventually there is an VNER part) *Let R be an execution of Algorithm 3.1 (which may include explicit delicate replacements). Eventually the system reaches a state (1) $c \in R$ after which an VNER part, R_{VNER}, starts. Moreover, after that, the system reaches eventually a state (2) $c_{\text{goodNtf}} \in R_{\text{VNER}}$: $\text{NA}(c_{\text{goodNtf}}) = \emptyset$ that has either (2.1) no notification or (2.2) at most one notification in the system and this notification becomes the system quorum eventually (after reaching a state in which invariants (1) to (7) of Claim 3.12 hold in c_{goodNtf}).*

Proof. Note that the proof of part (1) is done whenever R has a suffix R' in which no processor takes a step that includes a call to the function $estab$(set) (line 13). Let us suppose that suffix R' exists and show that part (2.1) is true. The reason is that R' satisfies the conditions of Lemma 3.8, which implies that eventually R' reaches c_{goodNtf}. This is because Lemma 3.8 says that the lexical largest notification is removed eventually from the system. Once that happen, either the system reaches c_{goodNtf}, or we can look at the suffix R'' that starts after that removal (of that lexical largest notification) and then we apply Lemma 3.8 again. Since R'' does not include steps with a call to the function $estab$(set), no new notifications are "added" to the system, and we can continue to apply Lemma 3.8 to the suffix of R'' until the system includes no notifications, i.e., it reaches c_{goodNtf}.

For the complementary case, we show that R includes an R_{VNER} part, i.e., part (1) of this proof. We consider the different cases of the phase of the notification with the maximal lexical value. We show that notification must reach the phase value of 2 and then leaves the system eventually. For this case, we show that parts (1) and (2.1) of this lemma hold.

Let us consider the case in which R includes an unbounded number of steps in which active processors call the function $estab$(set) (line 13). (Note that this is the complementary case to the first part of this proof because here R' does not exists.) Let $S = \{set : a_i \in R \text{ includes a step that calls } estab(\text{set})\}$. (Note that the case of $S = \emptyset$ is just another way to say that there are no effective steps in R and thus R' exists.) Suppose that $S \neq \emptyset$ and let $\text{n} = \langle \bullet, set \rangle$ be a notification with the largest lexical value in R, where set $\in S$. (Note that the proof needs not consider the case of $\text{n} = \langle 0, set \rangle$, see line 13.)

In order to complete the proof of parts (1) and (2) for the complementary case, we show that, for any choice of phase, it cannot be that $\text{n} = \langle phase, set \rangle$ and execution R does not have an VNER part, R_{VNER}. We do that by showing that the system reaches eventually a state c in which $noReco()$ (line 12) returns false for any active processor, $p_i \in P$. Moreover, c marks the beginning of R_{VNER} for which we show in part (2) of the proof that $c_{\text{goodNtf}} \in R_{\text{VNER}}$ follows.

Suppose that $\text{n} = \langle 1, set \rangle$**.** This part of the proof considers two cases. One in which n is a proposal that never leaves the system and the another case in which it does leave eventually.

• *Suppose that* n *never leaves the system.* Suppose, towards a contradiction, that (i) n never leaves the system and yet (ii) the system never reaches R_{VNER}. By the same arguments that appear in the proof of part (1) in Claim 3.12 we get that invariant $\text{prp}_i[j] = \text{n}$ holds, where $p_i, p_j \in P$ are any two active processors. Thus, $noReco()$ (line 12) returns false for any active processor, $p_i \in P$, and the system reaches R_{VNER}'s first configuration, c. Thus, a contradiction because the assumption that (i) and (ii) can hold together is false.

• *Suppose that* n *does leave the system.* We note that the case in which n leaves the system eventually contradicts our assumption that n has the highest lexical values in S. The reason is that only active processors can make a notification leave the system and they do so by changing their values in lines 28 to 28 (Claim 3.10). But then, for the case of $\text{n} = \langle 1, set \rangle$, the lexical value of n increases when Algorithm 3.1 takes such a step.

We note that there is no need to show part (2) of this lemma for the above two cases, because none of them is possible.

Suppose that n $= \langle 2, set \rangle$**.** Similar to the previous case, we consider both the case in which n is a proposal that never leaves the system and the other case in which it does leave eventually.

• *Suppose that* n *never leaves the system.* As in the proof of the first case of n $= \langle 1, set \rangle$ and in which n never leaves the system, we can use the same arguments as in the proof of part (1) of Claim 3.12 for showing that invariant $\text{prp}_i[j] = \text{n}$ holds, where $p_i, p_j \in P$ are any two active processors. Thus, $noReco()$ (line 12) returns false for any active processor, $p_i \in P$, and the system reaches R_{VNER}'s first configuration, c. Note that once the system enters the VNER part of the execution, R_{VNER}, we can use similar arguments to the ones used in the proof of Lemma 3.8 to show that the system removes n. Thus, a contradiction with our assumption (of this case) that n never leaves the system. Moreover, there is no need to prove part (2) of this lemma (for this impossible case).

• *Suppose that* n *does leave the system.* We show that the only possible case is that n $= \langle 2, set \rangle$ does leave the system. Let us consider the set $S' = \text{NA}(c_{\text{start}})$ of notifications that were present in R's starting state, c_{start}. We note that $|S'| \in \mathcal{O}(cap \cdot n^2)$, where cap is the bound on the link capacity (Section 2). Thus, even though the number of notifications in S is unbounded, the number of the notifications in S' is bounded. Since this case assumes that n eventually leaves the system, and all other cases above are impossible, it must be the case that R has a suffix R'' that starts in a system state $c'' : \text{NA}(c'') \cap S' \neq \emptyset$ the encodes no notification that appeared in R's starting state. This can be shown by induction on $|S'|$. Moreover, the proof can consider the case in which a given notification n $= \langle \bullet, \text{set} \rangle : \text{set} \in S'$ leaves the system and then returns via a call to $estab()$ (line 13). This is done by letting the proof to decorate the notification sets in R's starting system state, say, using the color 'red'. Moreover, each set in a notification that is the result of a call to $estab(\text{set})$ (line 13) also has decoration, but this time with another color, say 'blue'.

Let $\text{n}', \text{n}'' \in S \setminus S'$ be two notifications (with decorated sets) that become present in R but they are not present in R's starting state. The only way in which $\text{n}' = \langle 1, \bullet \rangle$ can become present in the system is via a step $a_1 \in R$ that calls $estab(\text{set})$ (line 13). Moreover, the only way that $\text{n}'' = \langle 2, \bullet \rangle$ can become present in R is via step $a_2 \in R$ that changes n′ to $\text{n}'' = \langle 2, set \rangle$ (lines 28 to 28), such that $\exists \text{n}' \in S \setminus S' : \text{n}'.set = \text{n}''.set$. Note that after step a_1, other steps $a_0 \in R$ occur in which n″ disappears from the system (lines 28 to 28). This sequence of steps between a_2 and a_0 and the system state that immediately precedes them are depicted by invariants (1) to (7) of Claim 3.12 since these invariants hold in these states. Thus, $c \in R$ holds immediately before a_2 and R must include an R_{VNER} part, because invariant (1) holds eventually in R. Moreover, invariants (1) to (7) of Claim 3.12 imply that n″ is the only notification in the system, which eventually becomes the quorum configuration after all steps a_0 are taken. Thus, part (2) of this lemma is shown. □

Theorem 3.15 demonstrates the eventual absence of stale information, which implies that Algorithm 3.1 convergences eventually, i.e., it is practically-stabilizing.

Theorem 3.15 (Convergence) *Let R be an admissible execution of Algorithm 3.1. Eventually the system reaches a state $c \in R$ in which the invariants of no type-1, type-2, type-3 and type-4 stale information hold thereafter.*

Proof. Claim 3.1 shows that there is no type-1 stale information eventually. Lemma 3.2 shows that there is no type-2 stale information eventually. Lemmas 3.8 and 3.14 say that the system either reaches a state in which there are no notifications or there is at most one notification (for which invariants (1) to (7) of Claim 3.12 hold) that later becomes the system configuration. Note that both cases imply that there is no type-3 stale information eventually. Claims 3.7 shows that there is no type-4 stale information eventually. □

Theorem 3.16 (Closure) *Let R be an execution of Algorithm 3.1. Suppose that execution R starts from a system state that includes no stale information. (1) for any system state $c \in R$, it holds that c includes no stale information. Suppose that the step that immediately follows c includes a call to $estab()$. (2) The only way that* set *becomes a notification is via a call to $estab$(set) (line 13) and the only way that a processor becomes a participant is via a call to $participate()$ (line 14). (3) The presence of notifications in R implies that one of them replaces the system configuration eventually.*

Proof. The proof essentially follows from established results above.
Part (1). Since there is no stale information in the system state that immediately proceeds c, there is no stale information in c. This follows from a close investigation of the lines that can change the system state in a way that might introduce stale information; the most relevant lines are the ones that deal with notifications (lines 28 and 28) and new participants (line 31). Thus, the proof is completed via parts (2) and (3).
Part (2). This is immediate from lines 13 and 14.
Part (3). It is not difficult to see that R includes an R_{VNER} part (Lemma 3.14, Part (1)). Then, the proof completes by applying part (7) of Claim 3.12 twice: Once for showing the selection of a single notification that has phase 2, and the second time for showing that the selected notification replaces the configuration. □

3.2 Reconfiguration Management

The Reconfiguration Management $recMA$ layer shown in Algorithm 3.2 bears the weight of initiating or *triggering* a reconfiguration when (i) the configuration majority has been lost, or (ii) when the prediction function $evalConf()$ indicates to a majority of members that a reconfiguration is needed to preserve the configuration. To trigger a reconfiguration, Algorithm 3.2 uses the $estab(set)$ interface with the $recSA$ layer. In this perspective, the two algorithms display their modularity as to their workings. Namely, $recMA$ controls *when* a reconfiguration should take place, but the reconfiguration replacement process is left to $recSA$, which will install a new configuration also trying to satisfy $recMA$'s proposal of the new configuration's set. Several processors may trigger reconfiguration simultaneously, but by the correctness of Algorithm 3.1 this does not affect the delicate reconfiguration, and by the correctness of Algorithm 3.2, each processor can only trigger once when this is needed.

In spite of using majorities, the algorithm is generalizable to other (more complex) quorum systems, while the prediction function $evalConf()$ can be either very simple, e.g., asking for reconfiguration once $1/4^{th}$ of the members are not trusted, or more complex, based on application criteria or network considerations. More elaborate methods may also be used to define the set of processors that Algorithm 3.2 proposes as the new configuration. Our current implementation, aiming at simplicity of presentation, defines the set of trusted participants of the proposer as the proposed set for the new configuration.

3.2.1 Algorithm Description

Preserving a majority. The algorithm strives to ensure that a majority of the configuration is active. Although majority is a special case of a quorum, the solution in extensible to host other quorum systems that can be built on top of the $config$ set, in which case, the algorithm aims at keeping a robust quorum system where robustness criteria are subject to the system's dynamics and application requirements. In this vein, the presented algorithm employs a configuration evaluation function $evalConf()$ used as a black box, which predicts the quality of the current $config$ and advises any participant whether a reconfiguration of $config$ needs to take place. Given that local information is possibly inaccurate, we prevent unilateral reconfiguration requests –that may be the result of inaccurate failure detection– by demanding that a processor must first be informed of a majority of processors in the current $config$ that also require a reconfiguration (lines 16–18).

Algorithm 3.2: Self-stabilizing Reconfiguration Management; code for processor p_i

1 **Interfaces:** $evalConf()$ returns True/False on whether a reconfiguration is required or not by using some (possibly application-based) prediction function. The rest of the interfaces are specified in Algorithm 3.1. $noReco()$ returns True if a reconfiguration is not taking place, or False otherwise. $estab(set)$ initiates the creation of a new configuration based on the set. $getConfig()$ returns the current local configuration.

2 **Variables:** $needReconf[\,]$ is an array of size at most N, composed of booleans $\{\mathsf{True}, \mathsf{False}\}$, where $needReconf_i[j]$ holds the last value of $needReconf_j[j]$ that p_i received from p_j as a result of exchange (lines 19 and 20) and $needReconf$ is an alias to $needReconf_i[i]$ i.e., of p_i's last reading of $evalConf()$. Similarly, $noMaj_i[\,]$ is an array of booleans of size at most N on whether some trusted processor of p_i detects a majority of members that are active per the reading of line 12. $noMaj_i[j]$ (for $i \neq j$) holds the last value of $noMaj_j[j]$ that p_i received from p_j. Finally $prevConfig$ holds p_i's believed previous $config$.

3 **Macros:**

4 $core() = \bigcap_{p_j \in FD_i[i].part} FD[j].part$;

5 $flushFlags()$: **foreach** $p_j \in FD[i]$ **do** $needReconf[j] \leftarrow (noMaj[j] \leftarrow \mathsf{False})$ **Do forever begin**

6 **if** $p_i \in FD[i].part$ **then**

7 $curConf = getConfig()$;

8 $needReconf[i] \leftarrow (noMaj[i] \leftarrow \mathsf{False})$;

9 **if** $prevConfig \notin \{curConf, \perp\}$ **then** $flushFlags()$

10 **if** $noReco() = \mathsf{True}$ **then**

11 $prevConfig \leftarrow curConf$;

12 **if** $|\{p_j \in curConf \cap FD[i]\}| < (\frac{|curConf|}{2} + 1)$ **then** $noMaj[i] \leftarrow \mathsf{True}$ **if** $(noMaj[i] = \mathsf{True}) \wedge (|core()| > 1) \wedge (\forall p_k \in core() : noMaj[k] = \mathsf{True})$ **then**

13 $estab(FD[i].part)$;

14 $flushFlags()$;

15

16 **else if** $(needReconf[i] \leftarrow evalConf(curConf)) \wedge |\{p_j \in curConf \cap FD[i] : needReconf[j] = \mathsf{True}\}| > \frac{|curConf|}{2}$ **then**

17 $estab(FD[i].part)$;

18 $flushFlags()$;

19 **foreach** $p_j \in FD[i].part$ **do** $send(\langle noMaj[i], needReconf[i]\rangle)$

20 **Upon receive** m **from** p_j **do if** $p_i \in FD[i].part$ **then** $\langle noMaj[j], needReconf[j]\rangle \leftarrow m$

Majority failure. On the other hand, we ensure liveness by handling the case were either the prediction function does not manage to sustain a majority, or an initial arbitrary state lacks a majority but there are no $config$ inconsistencies that can trigger a delicate reconfiguration (via the $estab()$ interface). Lines 12–14 tackle this case by defining the *core* of a processor p_i to be the intersection of the failure detector readings that p_i has for the processors in its own failure detector, i.e., $\cap_{p_j \in FD_i[i]} FD_i[j]$. If this local core agrees that there is no majority, i.e. that $noMaj = \mathsf{True}$, then p_i can request a new $config$. As a liveness condition to avoid triggering a new $config$ due to FD inconsistencies when there actually exists a majority of active configuration members, we place the *majority-supportive core* assumption on the failure detectors, as seen in Definition 3.2 below. Simply put, the assumption requires that if a majority of the current configuration is active, then the core of every processor p_t that is a participant, contains at least one processor p_s with a failure detector supporting that a majority of $config$ is trusted. Furthermore, p_t has knowledge that p_s can detect a majority of trusted members.

Detailed description. The algorithm is essentially executed only by participants as the condition of line 6 suggests. Line 7 reads the current configuration, while line 8 initiates the local $noMaj_i[i]$ and $needReconf_i[i]$ variables to False. If a change from the previous configuration has taken place, the arrays $noMaj[\,]$ and $needReconf[\,]$ are reset to False (line 9). The algorithm proceeds to evaluate whether

a reconfiguration is required by first checking whether a reconfiguration is already taking place (line 10) through the $noReco()$ interface of $recSA$. If this is not the case, then it checks whether it can see a trusted majority of configuration members, and updates the local $noMaj_i[i]$ boolean accordingly (line 12). If $noMaj_i[i] = \mathsf{True}$, i.e., no majority of members is active, and line 12 finds that all the processors in its core also fail to find a majority of members, then p_i can trigger reconfiguration using $estab()$ with the current local set of participants as the proposed new $config$ set (lines 13–14). The $needReconf_i[\,]$ and $noMaj_i[\,]$ arrays are again reset to False to prevent other processors that will receive these to trigger. Line 16 checks whether the prediction function $evalConfig()$ suggests a reconfiguration, and if a majority of members appears to agree on this, then the triggering proceeds as above. Participants continuously exchange their $noMaj$ and $needReconf$ variables (lines 19–20).

3.2.2 Correctness

The Reconfiguration Management algorithm is responsible for triggering a reconfiguration when either a majority of the members crash or whenever the (application-based) $config$ evaluation mechanism $evalConfig()$ suggests to a members' majority that a reconfiguration is required. The correctness proof ensures that, given the assumption of majority-supportive core holds, Algorithm 3.2 can converge from a transient initial state to a safe state, namely, that after $recMA$ has triggered a reconfiguration, it will never trigger a new one before the previous one is completed and only if a new event makes it necessary.

Terminology. We use the term *steady $config$ state* to indicate a system state in an execution were a $config$ has been installed by Algorithm 3.1 at least once, and the system state is conflict-free. A *legal execution R* for Algorithm 3.2, refers to an execution that converges to a steady $config$ state. Moreover, a reconfiguration in R takes place only when a majority of the configuration members fails, or when a majority of the members require a reconfiguration. The system remains conflict-free and moves to a new steady $config$ state with a new configuration.

Definition 3.2 (Majority-supportive core) *Consider a steady $config$ state in an execution R where the majority of members of the established $config$ never crashes. The* majority-supportive core *assumption requires that every participant p_i with a local core $\cap_{p_j \in FD_i[i]} FD_i[j]$ containing more than one processor, must have a core with at least one active participant p_r whose failure detector trusts a majority of the $config$, and for such a processor $noMaj_i[r] = \mathsf{False}$ throughout R.*

Remark 1 *We say that Algorithm 3.1 is* triggered *when a reconfiguration is initialized. By Algorithm 3.2, the only way that Algorithm 3.2 can cause a triggering of Algorithm 3.1 is through a call to the $estab()$ interface with Algorithm 3.1 on lines 13 and 17.*

Lemma 3.17 *Starting from an arbitrary initial state in an execution R, where stale information exists, Algorithm 3.2 converges to a steady $config$ state where local stale information are removed.*

Proof. Consider a processor p_i with an arbitrary initial local state where stale information exists (1) in the program counter, (2) in $noMaj_i[\bullet]$ and $needReconf_i[\bullet]$ and (3) in $prevConf$.
Case 1 – Stale information may initiate the algorithm in a line other than the first of the pseudocode. If p_i's program counter starts after line 10 and if a reconfiguration is taking place, then Algorithm 3.2 may force a second reconfiguration while Algorithm 3.1 is already reconfiguring. The counter for example could start on lines 13 and 17. This would force a brute reconfiguration. This triggering cannot be prevented in

such a transient state, but we note that any subsequent iteration of the algorithm is prevented from accessing $estab()$ lines (as in Remark 1) before the reconfiguration is finished.
Case 2 – We note that after a triggering as the one described above, the fields of arrays $noMaj_i[\bullet]$ and $needReconf_i[\bullet]$ are set to False. Moreover, in every iteration $noMaj_i[i]$ and $needReconf_i[i]$ are set to False. During a reconfiguration these values are propagated to other processors and p_i receives their corresponding values. Therefore, p_i must eventually receive $noMaj_j[j]$ ($needReconf_j[j]$) from some participant p_j, and overwrite any transient values. Lemma 3.18 bounds the number of reconfigurations that may be triggered by corruption that is not local, i.e., that emerges from corrupt $noMaj$ ($needReconf$) values that arrive from the communication links.
Case 3 – We anticipate that any reconfiguration returns a different configuration than the previous one. In a transient state though, the previous configuration ($prevConfig$) and the current configuration $curConf$ may coincide. This ignores the check of line 9 that sets $noMaj_i[\bullet]$ and $needReconf_i[\bullet]$ to False upon the detection of a reconfiguration change. This forms a source of a potential unneeded reconfiguration. Nevertheless, $prepConfig$ receives the most recent configuration value on every iteration of line 11 and thus the above may only take place once throughout R per processor.
Eliminating these sources of corruption, we reach a steady $config$ state without local stale information. □

Lemma 3.18 *Consider a steady $config$ state c in an execution R where the majority-supportive core assumption holds throughout, the majority of $config$ processors never crash and there is never a majority of members supporting $evalConf() =$ True. There is a bounded number of triggerings of the Algorithm 3.1 that are a result of stale information, namely $O(N^2 cap)$.*

Proof. By Remark 1, the only way that the algorithm may interrupt a steady $config$ state, is by reaching lines 13 and 17 that have a call to $estab()$. We assume that some member $p_t \in config$ has triggered Algorithm 3.1 at some system state $c_t \in R$, and we examine whether and when this state is attainable from c. We note that *in a complete iteration* of Algorithm 3.2, p_t must have no reconfiguration taking place while triggering, since this is a condition to reach the above mentioned lines imposed by line 10. We first prove that if there is a triggering it must be due to initial corrupt information and then argue that this can take place a bounded number of times.
Case 1 – The reconfiguration was initiated by line 13. This implies that the condition of line 12 is satisfied, i.e., at some system state $c_t \in R$, p_t has local information that satisfies $(noMaj_t[t] =$ True$)$ $\wedge$ $(|core_t()| > 1)$ $\wedge$ $(\forall p_k \in core_t() : noMaj_t[k] =$ True$)$. Condition $(noMaj_t[t] =$ True$)$ may be true locally for p_t, only due to failure detector inaccuracy, because, by the lemma, the majority of processors in the $config$ never fails throughout R. Condition $|core_t()| > 1$ suggests that p_t has at least two participant processors in its core (without requiring $p_t \in core_t()$). By the majority-supportive core assumption and the above, we are guaranteed that $\exists p_s \in core_t() : |FD_s[s] \cap config_s[s]| > \frac{|config_s[s]|}{2}$ throughout R and $noMaj_t[s] =$ True $\iff noMaj_s[s] =$ True. But in this case, $noMaj_s[s] =$ False and $noMaj_t[s] =$ True which contradicts the majority supportive assumption. We thus reach to the result.

Note that $noMaj_t[s] =$ True can reside in p_t's local state or in the communication links that may carry stale information. Because of the boundedness of our system, we can have one instance of corrupt $noMaj_t[s] =$ True in p_t's local state, and cap instances in the communication link. I.e., such information may cause a maximum of $1 + cap \cdot N$ triggerings per processor. Any processor that enters the system cannot introduce corrupt information to the system due to the data-links protocols and the joining mechanism. Thus majority supportive assumption is also attainable even when starting from arbitrary states.
Case 2 – The reconfiguration procedure was triggered by line 16. This implies that for p_t, both conditions were true, i.e., (a) $(needReconf_t \leftarrow evalConf_t(config_t))$ and (b) $|\{p_j \in config_t \cap FD_t :$

$needReconf_t[j] = \mathsf{True}\}| > \frac{|config_t|}{2}$. We note that the $needReconf_t[t]$ variable is always set to False upon the beginning of every iteration. Thus the local function $evalConf_t()$ due to p_t's failure detector and other application criteria explicitly suggested a reconfiguration *in the specific iteration* in which p_t triggered the reconfiguration. From the lemma, there is no majority of processors in the $config$ that supports a reconfiguration, even at the time when p_t triggered the reconfiguration.

Thus $needReconf_t[s] = \mathsf{True}$ must reside in p_t's local state and in the communication links. We can have one instance of corrupt $needReconf_t[s] = \mathsf{True}$ in p_t's local state, and cap instances in the communication links. I.e. such information may cause a maximum of $1 + cap \cdot N$ triggerings per processor. Note that after every such triggering, the source of triggering is eliminated by reseting $needReconfig[\,]$ to False (lines 9, 14 and 18). From this point onwards any processor that enters the system cannot by the data-links and the joining mechanism introduce corrupt information to the system.

So the possible triggerings in the system attributed to stale information are confined to $O(N^2 cap)$ and by Algorithm 3.1 guarantees we always reach a steady $config$ state. □

Let c_{safe} denote a *safe system state* where all possible sources of triggerings attributed to the arbitrary initial state have been eliminated. We denote an execution starting from c_{safe} as R_{safe}.

Lemma 3.19 *Consider an execution R_{safe} where the majority-supportive core assumption holds throughout, the majority of $config$ processors never crash and there is never a majority of the $config$ with local $evalConf() = \mathsf{True}$. This execution is composed of only steady $config$ states.*

Proof. By Lemma 3.18 there is a bounded number of triggerings due to initial arbitrary information. Given that we have reached a safe system state, these triggerings do not occur. The last $config$ change, has by line 9 reset all the fields in $noMaj[\,]$ and $needReconf[\,]$ to False and this holds for all participants (even if they are not members of the $config$). By our assumption a majority of processors does not crash. The majority-supportive core assumption states that throughout R_{safe} there exists at least one processor p_i in the core of p_t that always has $noMaj_i[i] = \mathsf{False}$ and p_t has $noMaj_t[i] = \mathsf{False}$. Thus the condition of line 12 can never be true, and thus there is no iteration of the algorithm that can reach line 13. Similarly, since no majority of processors in the $config$ change to $needReconf = \mathsf{True}$ in this execution, and the local states are exchanged continuously over the token-based data-link, line 17 cannot be true. Thus any system state in R is a steady $config$ state. □

Lemma 3.20 *Starting from an R_{safe} execution, Algorithm 3.2 guarantees that if (1) a majority of $config$ members collapse or if (2) a majority of members require a reconfiguration as per the prediction function, a reconfiguration takes place.*

Proof. We consider the two cases separately.
Case 1 – If a majority of the members collapses, then based on the failure detector's correctness, a non-crashed participant p_t will eventually stop including a majority of $config$ members in its failure detector and participants ($FD.part$) set. We remind that rejoins are not permitted. Since the majority-supporting core assumption does not apply in this case, any processor in p_t's core must eventually reach to the same conclusion as p_t. Every such participating processor $p_s \in core_t()$ propagates $noMaj_s[s] = \mathsf{True}$ in every iteration. By the assumption that a packet sent infinitely often arrives infinitely often (the fair communication assumption, Section 2), any processor such as p_t must eventually collect a $noMaj = \mathsf{True}$ from every member like p_s core and thus enable a reconfiguration.
Case 2 – The arguments are similar to Case 1. The difference lies in that the processor p_t must eventually receive $needReconfig = \mathsf{True}$ from a majority of $config$ members (rather than the local core processors) before it moves to trigger a reconfiguration. □

Lemma 3.21 *Starting from an R_{safe} execution, any triggering of Algorithm 3.1 (lines 13 and 17) related to a specific event (majority collapse or agreement of majority to change $config$), can only cause a one per participant concurrent trigger. After the $config$ has been established, no triggerings that relate to this event take place.*

Proof. We consider the two cases that can trigger a reconfiguration (Remark 1), and assume that p_t is the first processor to trigger $estab()$. Assume first that p_t has called Algorithm 3.1 two consecutive times, without a $config$ being *completely* established between the two calls. Note that a processor can access $estab()$ in either of lines 13 or 17 but not both in a single iteration. A call to $estab()$ initiates a reconfiguration and thus any subsequent check of p_t in line 10 returns False from p_t's $recSA$ layer. Thus p_t cannot access lines 13 or 17 until the reconfiguration has been completed. This implies that p_t can never trigger for a second time unless the new $config$ has been established. Note that if p_t triggers, another processor satisfying the conditions of line 10 may trigger concurrently, but is also subject to the trigger-once limitation. On the other hand, due to the exchange of information in Algorithm 3.1, when one processor triggers other processors eventually find their proposals and join the reconfiguration. So not every single processor's Reconfiguration Management module needs to trigger. Convergence to a single $config$ is guaranteed by Algorithm 3.1.

We conclude by indicating that lines 14 and 18 reset both arrays $noMaj_t[]$ and $needReconf_t[]$ immediately after $estab()$. Thus the triggering data used for this event are not used again. Moreover, upon configuration change, the same arrays are again set to False for the processors that have not triggered Algorithm 3.1 themselves through Algorithm 3.2. We thus reach a new steady $config$ state, and no more triggerings can take place due to the same event that had caused the reconfiguration. □

Theorem 3.22 *Let R be an execution starting from an arbitrary system state. Algorithm 3.2 guarantees that R eventually reaches an execution suffix which is a legal execution.*

Proof. By Lemmas 3.17 and 3.18, we are guaranteed that we reach a safe system state c_{safe} where stale information from the arbitrary initial state cannot force a triggering of new $config$. This is the suffix R_{safe}. Lemma 3.19 ensures that after we have reached c_{safe}, and until a new triggering takes place that is caused by a loss of majority or a majority of the $config$ deciding to reconfigure, the current $config$ will not be changed for any other reason. Lemma 3.21 guarantees that after a change, we return to a steady $config$ state. Hence, R_{safe} is a legal execution. □

3.3 Joining Mechanism

Every processor that wants to become a participant, uses the snap stabilizing data-link protocol (see Section 2) so as to avoid introducing stale information after it establishes a connection with the system's processors. Algorithm 3.1 enables a joiner to obtain the agreed $config$ when no reconfiguration is taking place. Note that, in spite of having knowledge of this $config$, a processor should only be able to participate in the computation if the application allows it. In order to sustain the self-stabilization property, it is also important that a new processor initializes its application-related local variables to either default values or to the latest values that a majority of the configuration members suggest. The joining protocol, Algorithm 3.3, illustrates the above and introduces joiners to the system, but only as *participants* and not as $config$ *members*.

The critical difference between a participant and a joiner is that the first is allowed to send configuration information via the $recSA$ layer, whereas the latter may only receive.

3.3.1 Algorithm description

The algorithm is executed by non-participants and participants alike.

Algorithm 3.3: Self-stabilizing Joining Mechanism; code for processor p_i

1 **Interfaces.** The algorithm uses following interfaces from Algorithm 3.1. $noReco()$ returns True if a reconfiguration is not taking place. $participate()$ makes p_i a participant. $getConfig()$ returns the agreed configuration from Algorithm 3.1 or $\perp$ if reconfiguration is taking place. The $passQuery()$ interface to the application, returns a True/False in response to granting a permission to a joining processor.
2 **Variables.** $FD[]$ as defined in Algorithm 3.1. $state[]$ is array of containing application states, where $state[i]$ represents p_i's local variables and $state[j]$ the state that p_i most recently received by p_j. $pass[]$ collects all the passes that p_i receives from configuration members.
3 **Functions.** $resetVars()$ initializes all variables related to the application based on default values. $initVars()$ initializes all variables related to the application based on states exchanged with the configuration members.
4 **procedure** $join()$ **begin**
5 **foreach** $p_j \in FD$ **do** $pass[j] \leftarrow$ False **do forever begin**
6 **if** $p_i \notin FD[i].part$ **then**
7 $resetVars()$;
8 **repeat**
9 **let** $comConf = getConfig()$;
10 **if** $noReco() \wedge (|\{p_j : p_j \in comConf \cap FD[i] \wedge pass[j] = \text{True}\}| > \frac{|comConf|}{2})$ **then**
11 $initVars()$;
12 $participate()$;
13 **foreach** $p_j \in FD[i]$ **do send**("Join")
14 **until** $p_i \in FD[i].part$;
15 **upon receive** ("Join") **from** $p_j \in FD \setminus FD[i].part$ **do begin**
16 **if** $p_i \in config \wedge noReco() =$ True **then send**$(\langle passQuery(), state_i \rangle)$
17 **upon receive** $m = \langle pass, state \rangle$ **from** $p_j \in FD$ **do begin**
18 **if** $p_i \notin FD[i].part$ **then** $\langle pass[j], state[j] \rangle \leftarrow m$

The joiner's side. Upon a call to the $join()$ procedure, a joiner sets all the entries of its $pass[\,]$ array to False (line 5) and resets application-related variables to default values, (lines 7). The processor then enters a do-forever loop, the contents of which it executes only while it is not a participant (line 6). A joiner then enters a loop in which it tries to gather enough support from a majority of configuration members. In every iteration, the joiner sends a "Join" request (line 13) and stores the True/False responses by any configuration member p_j in $pass[j]$, along with the latest application $state$ that p_j has send. If a majority of active members has granted a $pass =$ True and there is no reconfiguration taking place, then $participate()$ is called to allow the joining processor to become a participant.

The participant's side. A participant only executes the do–forever loop (line 5), but none of its contents since it always fails the condition of line 6. Participants however respond to join requests (line 16) by checking whether a joining processor has the correct configuration, and whether a reconfiguration is not taking place, as well as if the application can accept a new processor. If the above are satisfied then the participant sends a $pass =$ True and its applications' $state$, otherwise it responds with False.

3.3.2 Correctness

The term *legal join initiation* indicates a processor's attempt to become a participant by initiating Algorithm 3.3 on line 4, and not on any other line of the $join()$ procedure. If the latter case occurred it would indicate a corruption to the program counter.

Lemma 3.23 *Consider an arbitrary initial state in an execution R. There are up to N possible instances of processors introducing corruption to the system.*

Proof. Processors may be found with an uninitialized or falsely initialized local state due to a transient fault in their program counter which allowed them to reach line 11 without a legal join initiation. In an arbitrary initial state, any processor with stale information may manage to become a participant. There can be up to N such processors, i.e., the maximal number of active processors. Nevertheless, a processor trying to access the system after this, is forced to start the execution of $join()$ from line 4.

□

Claim 3.24 *Consider any processor p_i performing a legal join initiation. In the existence of other participants in the system, this processor never becomes a participant through the $join()$ procedure during reconfiguration.*

Proof. We consider the situation where participants exist and reconfiguration is taking place, thus $noReco()$ is False. In order for p_i to become a participant, it needs to gather a pass from at least a majority of the configuration members. This can only happen if a configuration is in place, and if each of these members is not reconfiguring. Thus if a pass is granted, it must be that during the execution more than a majority of True passes have arrived at p_i. Note that since the propagation of passes is continuous if a reconfiguration starts, then passes can also be retracted. Finally, since getting a majority of passes can coincide with the initialization of a reconfiguration, we note that due to asynchrony this processor is considered a participant of the previous configuration, since it has full knowledge of the system's state and is also known by the previous configuration members. □

Lemma 3.25 *Consider an execution R where Lemma 3.24 holds, such that during R, a processor p becomes a participant. Then p cannot cause a reconfiguration, unless there exists a majority of the configuration set, or if there is no majority of the $config$ that requires a reconfiguration.*

Proof. We assume that p enters the computation with a legal join initiation. If p triggers a reconfiguration in the absence of the above two cases, then this implies that p has managed to become a participant while carrying corrupt information which have triggered a reconfiguration either directly or indirectly (through Algorithm 3.2). Corruption can either be local or in the communication links. Since the snap-stabilizing data-link protocol runs before the processor calls $join()$, this removes data-link corruption for newly joining participants. We turn to the case of a corrupt local state. By the legal join initiation assumption, p must have reset its state on line 7. Before joining, the majority of members must acknowledge the latest state of p and p initiates its variables to legal values. It is therefore impossible that p can become a participant while it carries a corrupt state. Therefore, p cannot cause a reconfiguration. □

Theorem 3.26 *Consider an arbitrary initial state of an execution R of Algorithm 3.3. We eventually reach an execution suffix in which every joining processor p will continue trying to join a participant if the application allows it. Additionally, this new processor cannot trigger a delicate reconfiguration before becoming a participant and cannot trigger a delicate reconfiguration without majority loss or majority agreement after it becomes a participant.*

Proof. By Lemma 3.23, we eventually reach an execution suffix where all joining processors enter the computation with a legal join initiation. We assume that a reconfiguration is not taking place, that messages sent infinitely often are eventually received infinitely often, and that the application interface invoked by

the participating processors allows p to join. Then p will eventually have a failure detector including a majority of member processors and will send its "Join" request to a majority (line 13). Since there is no reconfiguration taking place, p must learn the current configuration from Algorithm 3.1, which should agree with the $config$ held by other processors. Thus each member must grant a pass to p by sending True through line 16. Therefore, p will gather a majority supporting its entrance and will eventually satisfy line 10. This allows it to reach line 12 and thus p becomes a participant. Notice that if the application does not give permission of entry via $passQuery()$, then p cannot become a participant unless this changes, but p will continue sending requests. Finally, Lemma 3.25 ensures that the new participant does not cause perturbations to the current configuration, and hence the result. □

4 Applications of the Reconfiguration Scheme

Using our self-stabilizing reconfiguration scheme, we present a collection of applications designed for more static settings (with a known fixed set of crash-prone processors) and adapt them to be able to run on the more dynamic setting that we describe here. We first present a general purpose labeling and counter scheme (Sections 4.1 and 4.2) and then proceed to show how to build a *self-stabilizing* reconfigurable virtually synchronous replicated state machine (Section 4.3).

4.1 Labeling Scheme and Algorithm

Many distributed applications assume access to an unbounded counter, e.g., to provide ballot numbers for consensus in Paxos, tag numbers for distributed shared memory emulation or view identifiers for virtually synchronous reliable multicast [11]. An unbounded counter implemented as a 64-bit integer, for example, is practically inexhaustible when initiated at 0 for the lifetime of most known systems. Transient failures, nevertheless, can immediately drive the counter (*sequence number* or $seqn$) to its maximal value (e.g., 2^{64}) causing it to wrap. Recently, we extended an existing labeling and counter increment scheme to enable any processor of a fixed processor set to increment a counter integer attached to an epoch $label$ [11]. When the counter is exhausted, a new maximal label is used with a non-exhausted $seqn$. We now adjust that solution to benefit from our reconfiguration mechanism. In this scheme, configuration members are the ones that run the labeling algorithm and maintain a globally maximal label and counter. The labeling and counter increment algorithms consider every new configuration as a new instance of the corresponding algorithms of [11]. To this end, Algorithm 4.1 is a wrapper of the self-stabilizing labeling algorithm of [11] that retains the initial algorithm as a module (Algorithm 4.2) allowing it to cope with reconfiguration. We first provide the labeling algorithm, and then extend labels to counters, presenting how counter increments take place.

4.1.1 Description of Algorithm 4.1

Outline. The algorithm is run only by configuration members. Each label is marked by its creator's identifier and any two labels are compared first as to their creator identifier and then as to a set of integers using the operator $\prec_{lb}$. Labels by the same processor can be *incomparable*. A processor that is aware of a set of labels with its own identifier, can always create a greater label. The aim is for members to learn of any valid label in the system and finally result to the globally greatest one, and to this end, members exchange labels. We refer the reader to [11] for more details on the label structure. The algorithm ignores labels by non-member creators by setting them to $\perp$. If Algorithm 3.1 reports that a reconfiguration is taking place (via $noReco$), no actions are taken. Upon the completion of a reconfiguration, every member's local label storage is rebuilt to reflect the new configuration set, and all the label queues are emptied. Newly joined

Algorithm 4.1: Self-Stabilizing Labeling Algorithm for Reconfiguration; code for $p_i \in config$

1 **Variables:** Let v be the size of the configuration $config$ as returned by $getConfig()$.
2 $max[v]$ of label pairs $\langle ml, cl\rangle$: $max[i]$ is p_i's largest label pair, $max[j]$ refers to p_j's $max_j[j]$ label pair that was last sent to p_i (canceled when $max[\bullet].cl \neq \bot$). $storedLabels[v]$: an array of queues of label pairs, where $storedLabels[j]$ holds the labels created by $p_j \in config$. For $p_j \in (config \setminus \{p_i\})$, $storedLabels[j]$'s queue size is limited to $(v + m)$ w.r.t. label pairs, where m is the maximum number of label pairs that can be in transit in the system. The $storedLabels[i]$'s queue size is limited to $(v(v^2 + m)) + v$ pairs.
3 **Interfaces:** $noReco()$ returns False (from the reconfiguration module) when a reconfiguration is taking place, and True otherwise. $getConfig()$ returns the current configuration if one exists. $labelReceiptAction()$ maintains the label arrays by calling the receive function of Algorithm 4.2.
4 **Operators:** $rebuild(v)$ rebuilds the $storedLabels[]$ array of queues and $max[]$ to have v entries. It also adjusts the queue size for the new v. $emptyAllQueues()$ clears all $storedLabels[]$ queues. $confChange()$ returns True if a reconfiguration has taken place, and False otherwise, by comparing the current label structures with the result of $getConfig()$.
5 **Macros:**
6 $cleanLP(x) =$ **if** $(\exists \ell \in x.\langle ml, cl\rangle, p_j \notin curConf) : \ell.lCreator = p_j)$ **then return** $\langle \bot, \bot\rangle$ **else return** x;
7 **function** $cleanMax()$ **foreach** $p_j \in curConf, \ell \in max[j].\langle ml, cl\rangle : \ell.lCreator = p_k \notin curConf$ **do** $max[j] \leftarrow \langle \bot, \bot\rangle$
8 **do forever begin**
9 **if** $noReco() =$ True $\wedge\ confChange() =$ True **then**
10 $curConf = getConfig()$;
11 $rebuild(|curConf|)$;
12 $emptyAllQueues()$;
13 $cleanMax()$;
14 $labelReceiptAction(\langle \bot, max[i], p_i\rangle)$;
15 **upon** $transmitReady(p_k \in curConf \setminus \{p_i\})$ **begin**
16 **if** $noReco() =$ True $\wedge\ confChange() =$ False **then**
17 **transmit**$(\langle max[i], max[k]\rangle \leftarrow \langle cleanLP(max[i]), cleanLP(max[k])\rangle)$
18 **upon** $receive(\langle sentMax, lastSent\rangle)$ **from** $p_k \in curConf$ **begin**
19 **if** $noReco() =$ True $\wedge\ confChange() =$ False **then**
20 $cleanMax()$;
21 $\langle sentMax, lastSent\rangle = \langle cleanLP(sentMax), cleanLP(lastSent)\rangle$;
22 $labelReceiptAction(\langle sentMax, lastSent, p_k\rangle)$;

processors are assumed to join with initialized links and empty label structures and thus cannot introduce corrupt information. If a reconfiguration is not reported, member p of the configuration periodically sends and receives its locally maximal labels with all the other members. Whenever it sends or receives a new label pair, it checks whether this has the identifier of one of the current members. The received label pairs are passed to the receive function (Algorithm 4.2), which is exactly the same as the one in [11]. This always returns a local maximal label either by some other member or by the caller itself. We underline that for every configuration we can find a greatest label, but it cannot be guaranteed that the label of a configuration will continue being the greatest in a following configuration.

Variables. Processor p_i that belongs to a configuration $config$ with $v = |config|$, has an array $max_i[v]$, where $max_i[i]$ contains the local maximal label p_i knows, and $max_i[j]$ the last value that $config$ member p_j has send. The array of label queues $storedLabels_i[]$ holds a queue of size $(v(v^2 + m)) + v$ in $storedLabels_i[i]$ for labels concerning processor p_i, and queues of size $v + m$ for all other configuration members in $storedLabels_i[j]$.

Algorithm 4.2: Self-Stabilizing Labeling Algorithm receipt action; code for $p_i \in config$

1 **Variables:** For a configuration $config$ with $v = |config|$
2 $max[v]$ of $\langle ml, cl\rangle$: $max[i]$ is p_i's largest label pair, $max[j]$ refers to p_j's label pair (canceled when $max[j].cl \neq \perp$).
3 $storedLabels[v]$: an array of queues of the most-recently-used label pairs, where $storedLabels[j]$ holds the labels created by $p_j \in config$. For $p_j \in (config \setminus \{p_i\})$, $storedLabels[j]$'s queue size is limited to $(v + m)$ w.r.t. label pairs, where m is the maximum number of label pairs that can be in transit in the system. The $storedLabels[i]$'s queue size is limited to $(v(v^2 + m)) + v$ pairs. The operator $add(\ell)$ adds lp to the front of the queue, and $emptyAllQueues()$ clears all $storedLabels[]$ queues. We use $lp.remove()$ for removing the record $lp \in storedLabels[]$. Note that an element is brought to the queue front every time this element is accessed in the queue.
4 $nextLabel()$ creates a label that is greater than any other label in $storedLabes_i[i]$.
5 **Notation:** Let y and y' be two records that include the field x. We denote $y =_x y' \equiv (y.x = y'.x)$
6 **Macros:**
7 $legit(lp) = (lp = \langle \bullet, \perp\rangle)$
8 $labels(lp)$: **return** $(storedLabels[lp.ml.lCreator])$
9 $double(j, lp) = (\exists lp' \in storedLabels[j] : ((lp \neq lp') \wedge ((lp =_{ml} lp') \vee (legit(lp) \wedge legit(lp')))))$
10 $staleInfo() = (\exists p_j \in P, lp \in storedLabels[j] : (lp \neq_{lCreator} j) \vee double(j, lp))$
11 $recordDoesntExist(j) = (\langle max[j].ml, \bullet\rangle \notin labels(max[j]))$
12 $notgeq(j, lp) =$ **if** $(\exists lp' \in storedLabels[j] : (lp'.ml \not\prec_{lb} lp.ml))$ **then return**$(lp'.ml)$ **else return**$(\perp)$
13 $canceled(lp) =$ **if** $(\exists lp' \in labels(lp) : ((lp' =_{ml} lp) \wedge \neg legit(lp')))$ **then return**(lp') **else return**$(\perp)$
14 $needsUpdate(j) = (\neg legit(max[j]) \wedge \langle max[j].ml, \perp\rangle \in labels(max[j]))$
15 $legitLabels() = \{max[j].ml : \exists p_j \in P \wedge legit(max[j])\}$
16 $useOwnLabel() =$ **if** $(\exists lp \in storedLabels[i] : legit(lp))$ **then** $max[i] \leftarrow lp$
else $storedLabels[i].add(max[i] \leftarrow \langle nextLabel(), \perp\rangle)$ `// For every` $lp \in storedLabels[i]$`, we pass in` $nextLabel()$ `both` $lp.ml$ `and` $lp.cl$`.`
17 **function** $labelReceiptAction(\langle sentMax, lastSent, p_k\rangle)$ **begin**
18 $max[k] \leftarrow sentMax$;
19 **if** $\neg legit(lastSent) \wedge max[i] =_{ml} lastSent$ **then** $max[i] \leftarrow lastSent$
20 **if** $staleInfo()$ **then** $storedLabels.emptyAllQueues()$
21 **foreach** $p_j \in P : recordDoesntExist(j)$ **do** $labels(max[j]).add(max[j])$
22 **foreach** $p_j \in P, lp \in storedLabels[j] : (legit(lp) \wedge (notgeq(j, lp) \neq \perp))$ **do** $lp.cl \leftarrow notgeq(j, lp)$
23 **foreach** $p_j \in P, lp \in labels(max[j]) : (\neg legit(max[j]) \wedge (max[j] =_{ml} lp) \wedge legit(lp))$ **do** $lp \leftarrow max[j]$
24 **foreach** $p_j \in P, lp \in storedLabels[j] : double(j, lp)$ **do** $lp.remove()$
25 **foreach** $p_j \in P : (legit(max[j]) \wedge (canceled(max[j]) \neq \perp))$ **do** $max[j] \leftarrow canceled(max[j])$
26 **if** $legitLabels() \neq \emptyset$ **then** $max[i] \leftarrow \langle \max_{\prec_{lb}}(legitLabels()), \perp\rangle$
27 **else** $useOwnLabel()$

Interfaces and Operators. The $noReco()$ interface of Algorithm 3.1 returns False when a reconfiguration is taking place, and True otherwise. It also returns the current configuration if one exists via $getConfig()$. $labelReceiptAction()$ maintains the label arrays by calling the receive function of Algorithm 4.2. $emptyAllQueues()$ clears all $storedLabels[]$ queues. $confChange()$ returns True if a reconfiguration has taken place, and False otherwise, by comparing the current label structures with the result of $getConfig()$.

During reconfiguration. The conditions $noReco() =$ True and $confChange() =$ False of lines 18 and 17 prevent transition and reception of labels during reconfiguration and before the label structures have been reset after reconfiguration has taken place.

After reconfiguration. Lines 9–14 are only executed upon a completed reconfiguration. Line 10 gets the new configuration from Algorithm 3.1, and line 11 uses $rebuild(config)$ to adjust $max[]$ array so that it holds the entries of any processor that also belonged to the previous configuration, removing the ones by old removed members and adding new fields for the labels of new $config$ members. Line 13 removes labels from the new $max[]$ that were not created by configuration members. The effect is analogous for $storedLabels[]$, where $v+m$-sized label queues are added and removed to reflect the changes in the $config$ composition, but noting that v is now the cardinality of the new configuration set. These queues are emptied by line 12. Finally, the processor finds a new local maximal label either from the ones that it has in $max[]$, or by creating a new one with its one creator identifier (line 14).

Label exchange and maintenance. If a reconfiguration is not taking place, then a member periodically sends to every other configuration member p_j its own local maximal label and the last label that p_j sent for the local maximal label. Note that messages from non-members are discarded before being sent, so they are not propagated (lines 15–17).

Similarly when p_i receives a message as the one described above, it cleans its $max[]$ array and the received two labels from non-member labels and passes them to the $labelReceptionAction()$ (Algorithm 4.2). The description of the inner workings of the $labelReceptionAction()$ is given in great detail in [11]. As an overview, it first stores p_j's value in $max_i[j]$ and processes it along with the other label that was received which reflected p_i's maximal label that was received most recently by p_j. In general terms, it performs housekeeping of labels in $storedLabels[]$ and $max[]$ such that only the greatest label per processor is considered before a local maximal label is chosen by p_i.

4.1.2 Correctness

Outline. We first establish that after a full iteration, of Algorithm 4.2 every configuration member does not sustain and does not introduce any label that has a creator identifier by a non-member. It is then possible to map every configuration of an execution to an instance of the fixed-set labeling algorithm of [11] and thus induce the correctness proof therein. This ensures that eventually a maximal label is found. Since the algorithm is aware of configuration changes (that exist after reconfigurations), it can use this events to achieve more efficient convergence for a configuration that follows a reconfiguration (in contrast to one that is in place in an initial arbitrary state). We thus reach to the above bounds where N is an upper bound on the system and thus a possible upper bound for the configuration size.

Note about participants that are not members. Algorithms 4.1 and 4.2 are run by processors that are strictly *members* of the current configuration and only if a reconfiguration is not taking place. Members send, receive and take into account only labels that come from and concern processors that are members. In this perspective, it should not be possible for a non-member processor to add labels to the system in a way that will affect the system. If a processor for any reason stops being a member after a configuration but remains an active participant, then once reconfiguration takes place, the member processors stop considering labels from this processor and *void* (set to $\langle \perp, \perp \rangle$) any label pair from this processor.

Lemma 4.1 *Consider an arbitrary starting system state of an execution R where a configuration does not change throughout R, a processor $p_i \in config$ reaches a local state where there are no labels from any processor $p_s \notin config$, immediately after the first iteration of Algorithm 4.1 that includes a receive event. Moreover, a label by $p_s \notin config$ can never be introduced to the local state of p_i at any point after the configuration is established.*

Proof. Notice that while the configuration does not change because the reconfiguration module finds that this configuration can serve the system, there could be corruption relating to the local variables of the labeling algorithm. We first establish that at any point after a complete execution of lines 18–22, p_i's local state never contains a label by a non-member p_s. Assume for contradiction that such a label exists at some system state after the execution of these lines. This label can exist either in (i) the $max[]$ array or (ii) the $storedLabels[]$ array.

Member $p_i \in config$ acknowledges that $p_s \notin config$, since the configuration is agreed and any inconsistency will cause a reconfiguration in Algorithm 3.1. We have already assumed for the purposes of the proof that a reconfiguration does not happen throughout R. Labels arriving from p_s cannot start a receive event by the conditions line 18. If a label created by p_s is received from some processor $p_j \in config$ then this is set to $\perp$ by line 21. So no incoming labels can enter the local state.

We now consider the system state immediately after line 22 returns from executing Algorithm 4.2. Line 20 of Algorithm 4.1 guarantees that $max[]$ is cleaned of labels that appear as created by p_s. Assume that a label by $p_s \notin config$ exists in the $storedLabels[]$ structure. There are two cases depending on the structure of the $storedLabels[]$ array of queues.

Case 1 – Processor p_i does not have an entry for p_s in $storedLabels[]$. This must be true for a legal state, since a member processor should only keep queues in $storedLabels_i[]$ that relate to member processors. So p_s's label must reside in a queue of $storedLabels_i[]$ that is not intended for p_s's labels. In such a case since $p_s \notin config$, line 20 of Algorithm 4.2 will cause the flushing of the queues because $staleInfo() = \mathsf{True}$.

Case 2 – Processor p_i has an entry for $p_s \notin config$ in $storedLabels[]$. This is the result of transient fault and implies that the label structures queues where not prepared for the new configuration. Nevertheless, in a complete iteration of lines 8–20, $confChange()$ of line 9 will return True because of this discrepancy between $config$ and the composition of the label structures. But this causes the algorithm to move to line 9 and so execute lines 10–13 thus emptying the $storedLabels[]$ and cleaning $max[]$ of p_s-created labels.

Therefore, immediately after a receive is completed for Case 1 or after the execution of line 13 in Case 2, there cannot be a label created by $p_s \notin config$ inside p_i's local state. Furthermore, we have established that such incoming labels cannot enter p_i's state. Hence we reach to the result. It is evident that once p_i's state is cleaned, it cannot transmit any such labels via line 17. □

Lemma 4.2 *Consider an execution R of Algorithm 4.1 where Lemma 4.1 holds and no reconfiguration takes place throughout R. It holds that this instance of the problem of providing a self-stabilizing labeling scheme in the dynamic setting can be reduced to the one of the problem of a fixed processor set self-stabilizing labeling scheme problem, where the fixed set is the common configuration.*

Proof. Note that the fixed-set version allows for processors from a specific non-changing set to crash but not rejoin. We identify the agreed configuration set to this fixed set of possibly active processors. Processors of the configuration may crash but may not rejoin. Algorithms 4.1 and 4.2 are run only by member processors. As established by Lemma 4.1, labels can have the identifier of any of the processors in this set but of no other processor. The communication links between configuration members are of bounded capacity and no labels from a non-member can be added to the local state of the member processors (since they are not considered by line 18) nor can be added to the communication links between the processors (line 17). Hence, the algorithm reduces every execution R to an instance of providing self-stabilizing labels. □

The corollary follows since we can use the solution to the fixed processor set problem to solve each instance of the problem whenever a reconfiguration is not taking place. In particular Algorithm 4.2 and line 17 comprise of the solution given in [11] to solve the fixed-set problem.

Corollary 4.3 *Consider an execution R of Algorithm 4.1 starting in an arbitrary state. While a reconfiguration is not taking place, the solution provided by [11] can be used to guarantee that a maximal label created by a member of the configuration will eventually be adopted by all active members.*

The theorem follows.

Theorem 4.4 *Starting in an arbitrary state, Algorithm 4.1 provides a maximal label. If a reconfiguration does not take place then there can be up to $O(N(N^2 + m))$ label creations before a maximal label is established. If a reconfiguration takes place then there can be up to $O(N^2)$ label creations.*

Proof. By Corollary 4.3 Algorithm 4.1 reaches a maximal label. By the results of Dolev et al [11] the worse case takes place when starting in an arbitrary state which requires at most $O(N(N^2 + m))$ label creations, where m is the system's communication link capacity in labels and in our case $N \geq v$. A processor may create a label with its own identifier the label it considers as maximal is canceled and it knows of no other labels. Labels that are possibly unknown to some processor and may cancel its maximal label, are found either in the local state ($|max[\bullet]| \leq N$ and $|storedLabels[N]| \leq 2N(N^2 + m - 1)$), or in transit between configuration members. Note that in an arbitrary state, the configuration may happen to be valid and not require a change, although the labels in the system may be corrupt. This is a worse case scenario that requires possibly $O(N(N^2 + m))$ label creations until the maximal label is reached.

Nevertheless, if a reconfiguration takes place, then the snap stabilizing link will clear the m labels found in the communication links and any active participant will empty their local queues in $storedLabels[]$. The only source of labels that could possibly force a processor to create new labels in order to find a maximal, are the $max[]$ arrays. There are up to N such arrays of size at most N, and hence $O(N^2)$ possible label creations. □

4.2 Counter Increment Algorithms

Using the labeling algorithm, we implement a practically infinite self-stabilizing counter, inexhaustible for the lifetime of most known systems (e.g., with upper bound of 2^{64}). We first show how to move from labels to counters and then provide a description of the algorithm along with a correctness proof.

4.2.1 Description

Counters. The labeling scheme used above, can be used to implement counters. The idea is to extend the labeling scheme to handle *counters*, where a counter is a triple $\langle label, seqn, wid \rangle$, where $seqn$ is an integer *sequence number*, ranging from 0 to 2^b, where b is large enough, say $b = 64$; and wid the processor identifier of the $seqn$ creator (not necessarily the same as the $label$'s creator). Specifically, we say that counter $ct_1 = \langle \ell_1, seqn_1, wid_1 \rangle$ is *smaller* than counter $ct_2 = \langle \ell_2, seqn_2, wid_2 \rangle$, and write that $ct_1 \prec_{ct} ct_2$, if ($\ell_1 \prec \ell_2$), or (($\ell_1 = \ell_2$) and ($seqn_1 < seqn_2$)), or (($\ell_1 = \ell_2$) and ($seqn_1 = seqn_2$) and ($wid_1 < wid_2$)). Note that when processors have the same label, the above relation forms a total ordering and processors can increment a shared counter also when attempting to do so concurrently. Also, when the labels of the two counters from the same processor are incomparable, the counters are also incomparable.

Outline. Algorithm 4.3 maintains counters as Algorithm 4.1 maintains labels. Counter increment for a participant that is not a configuration member is seen in Algorithm 4.5, and for a member, which also bears the responsibility to maintain the maximal counter, in Algorithm 4.4. A participant that wants to increment the counter, first queries the configuration for the maximal counter. It only needs to consider the responses

Algorithm 4.3: Self-stabilizing Counter Management Algorithm for Reconfiguration; code for $p_i \in config$

(Let v be the size of the configuration $config$ as returned by $getConfig()$).

Variables: A counter is a triple $\langle lbl, seqn, wid\rangle$ where lbl is a label of the labeling scheme, $seqn$ is the sequence number related to lbl, and wid is the identifier of the creator of this $seqn$. A counter pair $\langle mct, cct\rangle$ extends a label pair. cct is a canceling counter for mct, such that $cct.lbl \not\prec_{lb} mct.lbl$ or $cct.lbl = \bot$. We rename structures $max[]$ and $storedLabels[]$ of Alg. 4.2 to $maxC[]$ and $storedCnts[]$ that hold counter pairs instead of label pairs. Array $maxC[v]$ of counter pairs $\langle mct, cct\rangle$: $maxC[i]$ is p_i's largest counter pair, $maxC[j]$ refers to p_j's $maxC_j[j]$ counter pair that was last sent to p_i (canceled when $maxC[\bullet].cct.lbl \neq \bot$). $storedCnts[v]$: an array of queues of the counter pairs, where $storedCnts[j]$ holds the counters with labels created by $p_j \in config$. For $p_j \in (config \setminus \{p_i\})$, $storedCnts[j]$'s queue size is limited to $(v + m)$ w.r.t. counter pairs, where m is the maximum number of counter pairs that can be in transit in the system. The $storedCnts[i]$'s queue size is limited to $(v(v^2 + m)) + v$ pairs.

Interfaces: $noReco()$ returns False (from the reconfiguration module) when a reconfiguration is taking place, and True otherwise. $getConfig()$ returns the current configuration if one exists. $counterReceiptAction((\langle cnt, cnt, id\rangle)$ - executes the function $counterReceiptAction((\langle lbl, lbl, id\rangle)$ of Algorithm 4.2 adjusted for counter structures and handling counters. For counter pairs with the same mct label, only the instance with the greatest counter w.r.t. $\prec_{ct}$ is retained. In case where one counter is canceled, we keep the canceled. For ease of presentation we assume that a counter with a label created by p_i in line 27 of Algorithm 4.2, is initiated with a $seqn = 0$ and $wid = i$. A call of $counterReceiptAction()$ (without arguments) essentially ignores lines 18 and 19 of Alg. 4.2.

Operators: $rebuild(v)$ rebuilds the $storedCnts[]$ array of queues and $maxC[]$ to have v entries. It also adjusts the queue size for the new v. $emptyAllQueues()$ clears all $storedCnts[]$ queues. $confChange()$ returns True if a reconfiguration has taken place, and False otherwise, by comparing the current counter structures with the result of $getConfig()$. $enqueue(ctp)$ - places a counter pair ctp at the front of a queue. If $ctp.mct.lbl$ already exists in the queue, it only maintains the instance with the greatest counter w.r.t. $\prec_{ct}$, placing it at the front of the queue. If one counter pair is canceled then the canceled copy is the one retained.

Notation: Let y and y' be two records that include the field x. We denote $y =_x y' \equiv (y.x = y'.x)$.

Macros:

$cleanCP(x) =$ **if** $(\exists \ell \in x.\langle mct, cct\rangle.lbl, p_j \notin curConf) : \ell.lCreator = p_j)$ **then return** $\langle \bot, \bot\rangle$ **else return** x;

$exhausted(ctp) = (ctp.mct.seqn \geq 2^{64})$

$legit(ctp) = (ctp.cct = \bot)$

$cancelExhausted(ctp) : ctp.cct \leftarrow ctp.mct$

$cancelExhaustedMaxC()$: **foreach** $p_j \in config,\ c \in maxC[j] : exhausted(c)$ **do** $cancelExhausted(maxC[j])$

function $cleanMax()$ **foreach** $p_j \in curConf, \ell \in maxC[j].\langle mct, cct\rangle.lbl : \ell.lCreator = p_k \notin curConf$ **do** $maxC[j] \leftarrow \langle \bot, \bot\rangle$

do forever begin

  if $noReco() =$ True $\wedge\ confChange() =$ True **then**

    $curConf = getConfig()$;

    $rebuild(|curConf|)$;

    $emptyAllQueues()$;

    $cleanMax()$;

    $counterReceiptAction(\langle \bot, maxC[i], p_i\rangle)$;

upon $transmitReady(p_k \in curConf \setminus \{p_i\})$ **begin**

  if $noReco() =$ True $\wedge\ confChange() =$ False **then**

    transmit$(\langle maxC[i], maxC[k]\rangle \leftarrow \langle cleanCP(maxC[i]), cleanCP(maxC[k])\rangle)$

upon $receive(\langle sentMax, lastSent\rangle)$ **from** $p_k \in curConf$ **begin**

  if $noReco() =$ True $\wedge\ confChange() =$ False **then**

    $cleanMax()$;

    $\langle sentMax, lastSent\rangle = \langle cleanCP(sentMax), cleanCP(lastSent)\rangle$;

    $counterReceiptAction(\langle sentMax, lastSent, p_k\rangle)$;

upon request to increment counter $inc()$ **begin**

  if $noReco() =$ True $\wedge\ confChange() =$ False **then return** $incrementCounter(getConfig())$

of the majority, because the intersection property of majorities guarantees at least one member that holds the most recent value of a completed counter increment. Having this maximal counter, $seqn$ is incremented and written back to the configuration, awaiting for acknowledgments from a majority. This is, in spirit, similar to the two-phase write operation of MWMR register implementations, focusing on the sequence number rather than on an associated value.

Configuration members have the extra task of retaining the maximal value and ensuring the convergence property. To that end, they exchange their maximal counter and update their counter structures in the same way as in the labeling algorithm. The maximal counter needs to have the maximal label held by the configuration members and the highest sequence number, breaking symmetry with the writer identifier. If this maximal sequence number is *exhausted*, members proceed to find a new maximal label, using the maximal sequence number known for this epoch label (possibly 0 if it is a newly created or unused label).

Variables. The counter management algorithm (Algorithm 4.3) uses the same structures and procedures as the labeling algorithm, but now with counters instead of labels. Specifically we name the $max[\bullet]$ array to $maxC[\bullet]$ and the $storedLabels\bullet]$ array of label queues to $storedCnts[\bullet]$ array of counter queues. Processors only hold one copy of a counter with the same label, namely the one with the highest $seqn$ (breaking ties with wid). Processors send and receive pairs of counters $\langle mct, cct\rangle$ where the first is the believed *maximal* counter, and the second a *cancelling* counter. While $cct = \perp$ then mct is not-cancelled and can be considered as a valid candidate for the local maximal counter. When a counter becomes exhausted, i.e., its $seqn$ exceeds 2^b, mct becomes canceled by assigning $cct \leftarrow mct$.

Interfaces and Operators. $counterReceiptAction(\langle lbl, lbl, id\rangle)$ is an interface to Algorithm 4.2 which now acts on and maintains counter pairs rather than label pairs. It concludes by naming a local maximal label. We do not present this algorithm again as it is essentially the same, less naming "counter" in place of "label". The rest of the interfaces and operators are detailed at the beginning the algorithms in which they are being used for the first time.

Maintaining a maximal counter. Algorithm 4.3 presents how Algorithm 4.1 is modified to maintain a largest counter (rather than just the largest label), by exchanging the local maximal counters with other configuration members. The conditions that prevent send/receive during reconfiguration are the same. Similarly, the above actions do not take place after reconfiguration, until the structures are rebuild and reset in exactly the same way as described by the labeling algorithm. In addition, the same conditions that prevent send/receive also prevent incrementing the counter (line 29).

Counter increment for configuration members. Algorithm 4.4 shows how configuration members increment the counter. A member p_i first sends a query to all other members requesting the counter that they consider as the global maximum and awaits for responses from a majority. These counters are gathered and passed to the counter structures (line 19). Using the $counterReceiptAction()$ the algorithm eventually finds the maximal epoch label and the maximal sequence number it knows for this label. In other words, it collects counters and finds the counter(s) with the largest global label; there can be more than one such counter, in which case it returns the one with the highest sequence number, breaking symmetry with the sequence number processor ids. Then it checks whether this maximal sequence number is *exhausted*. When this is the case, it proceeds to find a new maximal label until it finds one that is not exhausted and uses the maximal sequence number it knows for this epoch label. The processor then increments the sequence number by one, sets its identifier as the writer of the sequence number (line 13) and sends the new counter

Algorithm 4.4: Counter Increment for configuration member $p_i \in config$; code for p_i

```
1  Variables, Interfaces and Macros are found in Algorithm 4.3.
2  Operator: abort() aborts the procedure and returns ⊥.
3  Notation: Let y and y' be two records that include the field x. We denote y =_x y' ≡ (y.x = y'.x).
4  Macros:
5  retCntrQ(ct) : return (storedCnts[ct.lbl.lCreator])
6  legitCounters() = {maxC[j].mct : ∃p_j ∈ config ∧ legit(maxC[j])}
7  getMaxSeq() : return max_wid({max_seqn({ctp : ctp.mct ∈ legitCounters() ∧ maxC[i] =_mct.lbl ctp})})
8  procedure incrementCounter(config) begin
9      majRead();
10     repeat
11         findMaxCounter();
12     until legit(maxC[i]) ∧ ¬exhausted(maxC[i]);
13     let newCntr = ⟨maxC[i].mct.lbl, maxC[i].mct.seqn + 1, i⟩;
14     if majWrite(newCntr) then maxC[i] ← newCntr;
15     retCntrQ(maxC[i].mct).enqueue(maxC[i])return maxC[i];
16 procedure majRead() begin
17     foreach p_j ∈ config do send majMaxRead() while waiting for responses from majority of config do
18         upon receipt of m from p_j ∈ config do
19         if m = Abort then abort() else maxC[j] ← cleanCP(m)
20 upon request for majMaxRead() from p_j begin
21     if noReco() then
22         findMaxCounter();
23         send maxC_i[i] to p_j;
24     else send Abort
25 procedure findMaxCounter() begin
26     cancelExhaustedMaxC();
27     counterReceiptAction();
28     maxC[i] ← getMaxSeq();
29 procedure majWrite(maxC_i[i]) begin
30     foreach p_j ∈ config do send majMaxWrite(maxC_i[i]) wait for ACK from majority of config;
31     if Abort received then abort()
32 upon request for majMaxWrite(max^j) from p_j begin
33     if noReco() = True then
34         maxC_i[j] ← cleanCT(max_ct(max^j, maxC_i[j]));
35         if max^j =_lbl.lCreator i then storedCnts_i[i].enqueue(maxC[i])if exhausted(maxC_i[j]) then
           cancelExhausted(maxC_i[j])send ACK to p_j;
36     else send Abort
```

to all members, awaiting for acknowledgments from a majority. Note that read and write requests during reconfigurations are answered with Abort. When the processor requesting the read/write receives an Abort it immediately terminates the increment procedure returning $\perp$.

Counter increment for non-member participants. As per Algorithm 4.5, participants that do not belong to the configuration, request counters from a majority of configuration members. They request the greatest with respect to $\prec_{ct}$ counter that is non-exhausted and legit. They then increment and write this to a majority of the configuration. If at any point during read or write they receive an Abort, they stop the procedure and return $\perp$. The same happens if they do not manage to find a maximal counter after the read. They can expect though that because of the counter propagation and the correctness of the labeling algorithm, such a counter

Algorithm 4.5: Counter Increment for non-member participant $p_i \notin config$; code for p_i

```
1  procedure incrementCounter(config) begin
2      let counters = majRead();
3      if ∀ct ∈ counter, ∃ct' ∈ counter : (¬exhausted(ct')) ∧ (legit(ct')) ∧ (ct ≺_ct ct') then
4          let maxCounter = ct' else let maxCounter = ⊥
5      if maxCntr ≠ ⊥ then
6          let newCntr = ⟨maxCntr.lbl, maxCntr.seqn + 1, i⟩;
7          if majWrite(newCntr) then return maxCntr else return ⊥
8      else return ⊥
9  procedure majRead() begin
10     let counters = ∅;
11     foreach p_j ∈ config do send majMaxRead() while |counters| ≤ |config|/2 do
12         upon receipt of m from p_j ∈ config do if m = Abort then abort() else counters.add(m)
13     return counters;
14 procedure majWrite(maxCntr) begin
15     foreach p_j ∈ config do send majMaxWrite(maxCntr) if ACK from majority of config received then wait;
16     else if Abort received then abort();
```

will eventually appear.

4.2.2 Correctness

Lemma 4.5 *Starting from an arbitrary state, in an execution R of Algorithm 4.3, configuration members eventually converge to a global maximal label.*

Proof. We note that the result asks that configuration members reach to a global maximal label, and thus it is not assume that they all hold the same $seqn$ and wid corresponding to this label. We prove in a step-by-step fashion, following the labeling algorithm lemmas and explaining how the new algorithms allow do not affect the correctness.
(i) Initially we note that labels by non-members are in an identical way as in the labeling algorithm excluded and cannot be reintroduced to the system (Lemma 4.1). This is immediate for Algorithm 4.3. Also lines 19 and 34 of Algorithm 4.4 guarantee that counters read during reads or writes are set to $\perp$ if they have labels by processors not in the current configuration.
(ii) In the same way as the labeling algorithm here is reduced the algorithm of [11] we perform the same reduction on the counter algorithm, by running the algorithm on the configuration member set and having each instance of the counter algorithm correspond to the fixed-set case. (iii) By this we conclude that we reach to a global maximal label for all active members. The bounds given by Theorem 4.4 are still relevant, since as we have explained, counters with a corrupted $seqn$ and counters with a corrupt label are bounded by the same numbers. Thus a counter may be adopted but then be exhausted very quickly, because it was initialized near its maximum. Thus eventually any corrupt counters will be removed. □

Theorem 4.6 *Starting from an arbitrary state, Algorithms 4.3, 4.4 and 4.5 eventually establish a monotonically increasing counter within the configuration that they are being executed.*

Proof. We first note that Lemma 4.5 ensures that processors reach to a maximal label. We establish that any two calls to $incrementCounter()$ return a strictly ordered counter value. Consider a processor p_i performing a counter increment. Given that no reconfiguration takes place and thus no aborts take place

during counter reads or writes, then if a majority has received the global counter cnt then by the intersection property of majorities, p_i must receive at least one copy of cnt when it reads (line 16 for Alg. 4.4 and line 9 for Alg. 4.5). For Algorithm 4.5 if any of the counters collected is legit and not exhausted it is incremented, and written back to a majority of configuration members with the writer's wid. In the members' algorithm, a member can always find a maximal counter (with Algorithm 4.3, even if the majority did not manage to return a legit, non-exhausted counter. Note that members always hold the greatest counter and discard of the smaller one. This ensures that after the write completes, any subsequent call to $incrementCounter()$ will return at least one copy of this new greatest value. Concurrent calls to $incrementCounter()$ may return to two processors p_i and p_j the same global maximal counter. This will be incremented by both to the same $seqn$, but will will be written with different wid's, so assuming $i < j$, any subsequent comparison will give the counter by p_j as the maximal. Hence the counter algorithm establishes a monotonically increasing counter. Also note, that for non-members, the algorithm will return a greater counter than the last completed call at the time of the call, or will return a $\perp$, in cases where reconfiguration is taking place, or where it is impossible to find a maximal counter, since the labels have yet to converge (and therefore incomparable counters exist). □

4.3 Reconfigurable Virtually Synchronous State Machine Replication

The self-stabilizing reconfiguration service together with the self-stabilizing labeling and counter scheme, can extend the capabilities of various applications to run on more dynamic settings. Virtual synchrony is an established technique for achieving state machine replication (SMR). We have recently presented a self-stabilizing virtually synchronous SMR algorithm for a fixed set of processors [11]. We now discuss how this virtually synchronous SMR solution can benefit from the reconfiguration service to run on more dynamically changing environments. We note that applications can use the reconfiguration service to guarantee continuous service, only when a delicate reconfigurations take place between periods of steady configuration and not when brute reconfiguration takes place. Nevertheless, brute reconfiguration guarantees that, after a transient fault, the service will eventually return to the desired behavior.

A *view* is a the set of processors with a unique identifier, that allows to the view members to achieve reliable multicast within the view. The self-stabilizing reconfigurable virtually synchronous SMR task is specified such that any two processors that are together in any two consecutive views will deliver the same messages in their respective views and preserve the same state. Moreover, the two views can be belong to two different consecutive configurations when the second configuration was the result of delicate reconfiguration. A pseudocode is found in Algorithm 4.7 and a description and correctness proof follows.

4.3.1 Description

Obtaining a coordinator. The virtual synchrony algorithm of [11] is coordinator-based and works on the primary component given the *supportive majority* assumption on the failure detectors. This assumption states that a majority of processors of the (fixed) processor set mutually never suspect some processor on their failure detectors throughout an infinite execution, given that this processor does not crash. The proof of [11] establishes that such a supported processor eventually becomes the coordinator throughout the execution.

We modify the definition of supportive majority so that a majority of the configuration members' set ($config$) needs to provide such support to the coordinator. Out of the current $config$, any $config$ member with supporting majority can obtain a counter from the labeling and counter algorithms that are run on $config$. The processor with the greatest counter, becomes the coordinator and establishes a *view* reflecting

Algorithm 4.6: Coordinator-led delicate reconfiguration; code for processor p_i

1 **Interfaces:** $needDelicateReconf()$ returns True/False on whether p_i is the coordinator and whether application criteria requires and is ready for a a reconfiguration.

```
/* Replaces line 16 of Reconfiguration Management layer                    */
```

17 **else if** $needDelicateReconf()$ **then**

its own local failure detector, and with the counter as its view identifier[1]. In doing so it also collects the states from view members as well as undelivered messages and synchronizes to create the most up-to-date state for the view members to replicate. The coordinator changes the view when the view set does not reflect its failure detector set. View members follow the view composition and state of the coordinator. Upon coordinator collapse, the same process provides a new coordinator and preserves the state. Note that the correctness for this is immediate from the correctness proofs of [11].

Coordinator-controlled joins. By the joining protocol, before a joiner tries to become a participant, it has any application-related local data set to $\perp$. The coordinator is the one that controls whether the application may or may not allow joining processors to become participants. Several approaches may be used to give permission to joiners. The coordinator may allow joining whenever it detects that more participants are required, by raising a flag and warning $config$ members that they can allow access to the application. This can be easily applied to our current joining protocol (Algorithm 3.3) where configuration members implement the $passQuery()$ interface, simply by returning their most recent True/False value they have for the coordinator's flag. Another approach would be for a configuration member to provide the coordinator's details to the joiner so that joiners may, by directly communicating with the coordinator, gain permission to become participants.

Note that in both cases, joiners may become participants but they have yet to gain access to the application. This takes place if they are part of the coordinator's FD and are subsequently included in the next view. In becoming part of the view they also acquire a copy of the most recent state and begin state replication. This gives leverage to the application when controlling the number of processors that are allowed to access the application.

Coordinator-initiated reconfiguration. The coordinator can be given the authority to initiate delicate reconfigurations. This implies that there is no need for the prediction functions of a majority of processors to support a reconfiguration before it can take place, but it only suffices for a coordinator to be in place, in order to take the decision for reconfiguration based on application-specific criteria and suspend changes to the state during reconfiguration. To this end, line 16 of Algorithm 3.2 that triggers the delicate reconfiguration in Algorithm 3.1 is replaced with line 17 of Algorithm 4.6. The $needReconf$ flag is rendered unnecessary for this cause.

Before reconfiguration. In order to initiate a reconfiguration that will not result in loss of the state of the replicas, the coordinator must ensure that all view processors have the most recent state and that the exchange of messages is suspended. We require that for the state to survive to the first view of the next configuration, at least one replica of the last view before reconfiguration must not crash.

[1]We may occasionally refer to view members and this should not be confused with configuration members, although a majority of configuration members must belong to the view

Algorithm 4.7: Self-stabilizing reconfigurable VS SMR; code **for participant** p_i

Interfaces: $fetch()$ next multicast message, $apply(state, msg)$ applies the step msg to $state$ (while producing side effects), $synchState(replica)$ returns a replica consolidated state, $synchMsgs(replica)$ returns a consolidated array of last delivered messages, $inc()$ returns a counter from the increment counter algorithm, $getConfig()$ returns the latest configuration and $noReco()$ returns True/False on whether a reconfiguration is *not* taking place (Algorithm 3.1), $evalConf()$ is the reconfiguration prediction function returning True/False.

Variables: The following arrays consider both p_i's own value (the i-th entry) and p_j's most recently received value (the j-th entry). $state[\bullet] = \langle view = \langle ID, set\rangle, status \in \{\text{Multicast, Propose, Install}\}, (multicast\ round\ number)\ rnd, (replica)\ state, (last\ delivered\ messages)\ msg[n]\ (to\ the\ state\ machine), (last\ fetched)\ input\ (to\ the\ state\ machine), propV = \langle ID, set\rangle, (no\ coordinator\ alive)\ noCrd, (message\ delivery)\ suspend\rangle$: an array of state variables. $FD[j].crd$ returns the last reported identifier of p_j's local coordinator which is continuously propagated using the token exchange (see Section 2). It is $\perp$ if p_j has no coordinator.

Function $needDelicateReconf() = (reconfReady = \text{True} \wedge valCrd = p_i \wedge evalConfig() = \text{True})$;

Do forever begin

- **let** $curConf = getConfig()$; `// Gets latest config set`
- **let** $seemCrd = \{p_\ell = state[\ell].propV.ID.wid \in FD.part \cap curConfig : (|state[\ell].propV.set| > \lfloor |curConf|/2\rfloor)\ (|state[\ell].FD.part| > \lfloor config/2\rfloor) \wedge (p_\ell \in state[\ell].propV.set) \wedge (p_k \in state[\ell].propV.set \leftrightarrow p_\ell \in FD[k].part) \wedge ((state[\ell].status = \text{Multicast}) \rightarrow (state[\ell].(view = propV) \wedge crd(\ell) = \ell)) \wedge ((state[\ell].status = \text{Install}) \rightarrow FD[\ell].crd = \ell)\}$;
- **let** $valCrd = \{p_\ell \in seemCrd : (\forall p_k \in seemCrd : state[k].propV.ID \preceq_{ct} state[\ell].propV.ID)\}$;
- $noCrd \leftarrow (|valCrd| \neq 1)$; $FD[i].crd \leftarrow valCrd$;
- **if** $(valCrd = \{p_i\} \wedge status = \text{Multicast} \wedge reconfReady = \text{True})$ **then** $suspend \leftarrow (reconfReady \leftarrow evalConfig())$ **else if** $(valCrd \neq \{p_i\}) \wedge (valCrd = \{p_\ell\} \wedge state[\ell].status \in \{\text{Propose, Install}\})$ **then** $suspend \leftarrow (reconfReady \leftarrow \text{False})$ **if** $noReco() = \text{False}$ **then** $suspend \leftarrow \text{True}$
- **if** $(|FD.part \cap curConf| > \lfloor |curConf|/2\rfloor) \wedge (((|valCrd| \neq 1) \wedge (|\{p_k \in FD.part : p_i \in FD[k].part \wedge state[k].noCrd\}| > \lfloor |curConf|/2\rfloor)) \vee ((valCrd = \{p_i\}) \wedge (FD.part \neq propV.set) \wedge (|\{p_k \in FD.part : state[k].propV = propV\}| > \lfloor |curConf|/2\rfloor))) \wedge noReco()$ **then** $(status, propV) \leftarrow (\text{Propose}, \langle inc(), FD.part\rangle)$
- **else if** $(valCrd = \{p_i\}) \wedge (\forall\ p_j \in view.set : state[j].(view, status, rnd) = (view, status, rnd)) \vee ((status \neq \text{Multicast}) \wedge (\forall\ p_j \in propV.set : state[j].(propV, status) = (propV, \text{Propose}))$ **then**
  - **if** $status = \text{Multicast} \wedge reconfReady = \text{False}$ **then**
    - $apply(state, msg)$; $suspend \leftarrow evalConf()$;
    - **if** $(reconfReady \leftarrow (\forall p_k \in view.set : rep[k].suspend = \text{True})) = \text{False}) \vee (\text{noReco}() = \text{True})$ **then**
      - $input \leftarrow fetch()$;
      - **foreach** $p_j \in P$ **do if** $p_j \in view.set$ **then** $msg[j] \leftarrow state[j].input$ **else** $msg[j] \leftarrow \perp$
      - $rnd \leftarrow rnd + 1$;
  - **else if** $status = \text{Propose}$ **then** $(state, status, msg) \leftarrow (synchState(state), \text{Install}, synchMsgs(state))$ **else if** $status = \text{Install}$ **then** $(view, status, rnd, suspend, reconfReady) \leftarrow (propV, \text{Multicast}, 0, \text{False}, \text{False})$
- **else if** $valCrd = \{p_\ell\} \wedge \ell \neq i \wedge ((state[\ell].rnd = 0 \vee rnd < state[\ell].rnd \vee state[\ell].(view \neq propV))$ **then**
  - **if** $state[\ell].status = \text{Multicast} \wedge suspend = \text{False}$ **then**
    - $state[i] \leftarrow state[\ell]$; `/* also adopts suspend flag */`
    - $apply(state, state[\ell].msg)$; `/* for the sake of side-effects */`
    - **if** $state[\ell].suspend = \text{False}$ **then** $input \leftarrow fetch()$
  - **else if** $state[\ell].status = \text{Install}$ **then** $state[i] \leftarrow state[\ell]$ **else if** $state[\ell].status = \text{Propose}$ **then** $(status, propV) \leftarrow state[\ell].(status, propV)$
- **let** $sendSet = (seemCrd \cup \{p_k \in propV.set : valCrd = \{p_i\}\} \cup \{p_k \in FD : noCrd \vee (status = \text{Propose})\})$
- **foreach** $p_j \in sendSet$ **do** $send(state[i])$

Upon message arrival m **from** p_j **do** $state[j] \leftarrow m$;

When the coordinator is informed by its prediction function that reconfiguration is required, it performs a multicast round to gather the most recent state and received messages, but also raises a $suspend$ flag that

it propagates with the current state. It waits until all view members return their current state and messages and that they have suspended new messages. Upon receiving this information it performs a final multicast round to ensure that the last state with the messages have been received and applied and then calls the reconfiguration through Algorithm 3.2. Note that in case of coordinator crash, a new view coordinator needs to be established before the reconfiguration can take place, which adds an extra delay.

After reconfiguration. Once the labeling and increment counter algorithm have stabilized to a new max counter any member of the new configuration has access to a counter value in order to become the coordinator. Processors propagate their state which can be either the last before reconfiguration or $\sharp$ in case they are newly joining processors. The new coordinator will synchronize the states and messages and establish the new view so that the service can continue, and will also resume the application for new messages to be fetched. This is in essence a mere view installation with nothing additional to the algorithm of [11].

Providing service. Inside the view, the basic functionality of the algorithm of [11], i.e., reliable multicast and state replication, are not in any way obstructed by the underlying reconfiguration service when a reconfiguration is not taking place. Additionally, view changes are not affected when reconfiguration is not taking place.

4.3.2 Correctness

We establish the correctness of the algorithm, by extending the correctness results of [11]. We first prove that the proposed service (Algorithm 4.7), when working in an established configuration, is an execution of the algorithm of [11] and thus the correctness for that case follows. We then consider the stabilization of the application after a reconfiguration, and conclude by establishing that state replication is unaffected as to its correctness when a delicate reconfiguration that was initiated by the coordinator takes place.

Lemma 4.7 *Starting in an arbitrary state in an execution R, stale information can only cause a single reconfiguration throughout R, unless the prediction mechanism $evalConfig()$ allows it.*

Proof. We note that stale information relating to reconfiguration consist of $suspend$ and the copies of $suspend$ exchanged with every other processor, and $reconfReady$, as well as $valCrd$ that may result from stale information as per lines 6 and 7. We pay attention to when the conditions are satisfied for the $needDelicateReconf()$ interface (line 3). This is called by a processor that believes itself to be the coordinator ($valCrd = p_i$), in order to initiate a delicate reconfiguration in the modified Algorithm 4.6. The three conditions are: *(i)* $evalConfig()$ returns True, *(ii)* All processors of the view have $suspend =$ True, *(iii)* The caller believes itself to be the coordinator, i.e., $valCrd = p_i$. We will refer to a coordinator as one who has established its view in the primary component of the $config$, and other $config$ members follow this processor in performing state replication. We study the following three complementary cases.
Case 1 – A non-coordinator p_i that does not believe itself to be the coordinator. (I.e., $valCrd \neq p_i$) In this case, p_i cannot initiate a reconfiguration since condition (iii) will always fail.
Case 2 – A non-coordinator p_i that believes it is the coordinator. (I.e., $valCrd = p_i$ locally for p_i, but $valCrd \neq p_i$ for a set of other processors greater than the majority of the $config$.) In [11] processors such as p_i are proved to eventually stop believing to be the coordinator due to propagation of information assumption. Nevertheless, before information may be propagated, such a processor's stale information, can cause a reconfiguration if the conditions are satisfied. Note that after reconfiguration takes place, this processor needs to again establish its (alleged) coordinatorship before it reconfigures again and in the process it needs

to reset its $suspend$ variables (line 9.

Case 3 – A coordinator p_i with corrupt initial state, cannot cause a reconfiguration in case $evalConfig() = \mathsf{False}$. It may cause a reconfiguration if also satisfies condition *(ii)* due to corruption and $evalConfig() = \mathsf{True}$. This cannot be distinguished from the non-transient case. As already noted, after reconfiguration takes place, this coordinator needs to again establish its coordinatorship and in the process it needs to reset its $suspend$ variables (line 9). □

Lemma 4.8 *Consider an execution R of Algorithm 4.7 in which the following hold throughout: (i) the supportive majority assumption, (ii) no reconfiguration takes place (and therefore a valid configuration $config$ is in place) and (iii) $evalConfig() = \mathsf{False}$. Then this execution is a reduction to the execution of the self-stabilizing virtually synchronous SMR algorithm of [11].*

Proof. The mapping of the fixed set of processors of [11] (that was named P where $|P| = n$) to the set of $config$, creates the fixed set (throughout R) that provides the supportive majority (of $\lfloor |config|/2 \rfloor$) for at least one processor of the configuration to eventually become the coordinator. While the views can also include non-member processors (that are participants), safety is provided by the majority of $config$ members. We note that the algorithm in its structure has not changed less the fact of replacing the old fixed set with $config$, and the addition of the suspend mechanism which is not activated at any point, since $evalConfig() = \mathsf{False}$ throughout R. We thus deduce that the lemma is a reduction to the non-reconfigurable VS SMR and we deduce the following corollary. □

Corollary 4.9 *Consider an execution R of Algorithm 4.7 in which the following hold throughout: (i) the supportive majority assumption, (ii) no reconfiguration takes place (and therefore a valid configuration $config$ is in place) and (iii) $evalConfig() = \mathsf{False}$. Then, by Lemma 4.8 and the correctness of [11], starting from an arbitrary state, Algorithm 4.7 simulates state machine replication preserving the virtual synchrony property.*

Lemma 4.10 *Consider an infinite execution R of Algorithm 4.7 with an established coordinator p_i. If $evalConfig_i() = \mathsf{True}$ from some system state $c \in R$ onwards, then we reach another system state $c' \in R$ in which a reconfiguration eventually take place (and after which $evalConfig() = \mathsf{False}$).*

Proof. From c onwards, the established coordinator p_i receives True whenever it calls $evalConfig()$. This implies that because of line 13 $state_i[i].suspend = \mathsf{True}$ always. This value is sent in every iteration of the Algorithm by p_i (lines 24–25) to every processor in the view, and adopted by every processor in the view through lines 26 and 20. Every view member p_j adopting p_i's state in $state_j[i]$ (along with $suspend = \mathsf{True}$) propagates $state_j[j].suspend = \mathsf{True}$ back to p_i. Due to assumption that a message sent infinitely often is received infinitely often, eventually the coordinator learns that every processor in its view has adopted $suspend = \mathsf{True}$.

If during this process the coordinator needs to change the view (due to failure detector changes), then this should take place and the above procedure starts again after the view installation (because installation by line 17 forces $suspend = \mathsf{False}$. The same takes place if the coordinator is lost, in which case by line 9 and the previous remark about installation falsify $suspend$ flags, and the procedure for reconfiguring should be taken up by the next coordinator.

If every processor in the view has $suspend = \mathsf{True}$ and p_i has been informed of this, then line 14 sets $reconfReady$ to True for the coordinator, and line 9 retains this value since $evalConfig() = \mathsf{True}$ by assumption. Hence, whenever the coordinator uses the reconfiguration manager (Algorithm 4.6) to run

line 17, it will find that all conditions of $needDelicateReconf()$ are satisfied, thus the coordinator can move on to initiate the reconfiguration using the $configEstb()$ interface of the reconfiguration scheme, and hence the result. □

Lemma 4.11 *Consider an infinite execution R of Algorithm 4.7, starting from a system state c' as in Lemma 4.10. We eventually reach a new state $c'' \in R$ in which a new configuration is installed and a new valid coordinator is established in the primary component of $config$.*

Proof. Starting from system state c', a coordinator p_i has already initiated a reconfiguration. During reconfiguration, p_i cannot proceed to increment the round number since line 14 allows no round increments if reconfiguration is taking place. Note that since it is the coordinator that triggers the reconfiguration, this processor is immediately informed about a reconfiguration through $noReco()$. By the correctness of the reconfiguration and stability assurance layer, we are guaranteed that the reconfiguration will complete and so we move from a configuration $config$ to a new one $config'$.

When the reconfiguration completes, every processor in the new configuration $config'$ should try to establish that there is a coordinator installed. Even if p_i survives to $config'$, it will need to change the view because of the change in configuration, although it will still consider itself as the coordinator. By the second set of conditions that allow view proposal, if $((valCrd = \{p_i\}) \wedge (FD_i[i].part \neq propV.set) \wedge (|\{p_k \in FD_i[i].part : state_i[k].propV = propV_i\}| > \lfloor |config'|/2 \rfloor)))$, namely if there is a majority of active maembers that follow the previously proposed view of p_i, then p_i is eligible to create a new view. Otherwise, any processor p_ℓ may trigger, given that the following conditions are satisfied. (i) $|FD_\ell[\ell].part| \cap config' > \lfloor |config'|/2 \rfloor$ (p_j can trust a majority of the configuration). (ii) The set of processors in $FD_\ell[\ell].part$ that p_ℓ trusts and p_ℓ knows to have their trust must state that they have no coordinator. Note that by our majority supportive assumption and the eventual reception of messages, some processor will reach a state in which it will be able to propose.

The correctness arguments of [11] complete the proof (by Corollary 4.9), and so some processor manages to become the coordinator. □

Lemma 4.12 *Consider an infinite execution R of Algorithm 4.7 that contains a system state c' as defined in Lemma 4.10 and a following state c'' as guaranteed by Lemma 4.11. The replica state is preserved from c' to c''.*

Proof. We examine how the state at the time of reconfiguration triggering, during and after reconfiguration remains unchanged until a coordinator is in place and initiates multicasting.

Step 1 – We establish that the state of the replicas **before** reconfiguration, is preserved. By Lemma 4.10, processor are led to turn their flag $suspend$ to False, when the coordinator p_i suggests that a reconfiguration must take place. In the next multicast round when every non-coordinator processor receives p_i's $suspend =$ True it adopts the last state (line 20), applies changes to the state 21 and fetches multicast messages *only if the suspend flag is* False. Similarly, the coordinator applies the changes to the state 21 after all the view members have adopted its state, and renews its suspend flag (which should always return True, otherwise the suspension process stops). If all the participants have suspended (otherwise condition ($\forall\ p_j \in view.set : state_i[j].(view, status, rnd) = state_i[i].(view, status, rnd)$) of line 10 would not hold), then line 14 must return True so $reconfReady =$ True and the coordinator does not fetch new messages (line 15) and stops incrementing round numbers.

Step 2 – During reconfiguration, no new multicasts and message $fetch()$ may take place, since non-coordinators of $config$ cannot access line 19 due to the conditions of 18. They also cannot propose a

new view during reconfiguration, by the last condition of line 10. So their $suspend$ flags are preserved to $sfTrue$ throughout the reconfiguration (also by line 9). The coordinator's flag is also maintained to True if reconfiguration is taking place.

Step 3 – After the new configuration $config'$ is in place, new processors are assumed to access the computation with default values in their state, and cannot therefore introduce new messages or replica state. Processors coming from $config$, carry the last state and for liveness we assume a single of these processors is alive and included in the proposed view set of the next coordinator. That a new proposal takes place, it is suggested by Lemma 4.11. So the new coordinator (possibly the same as the one before reconfiguration) will gather all available states and messages and create a synchronized state (which is only for stabilization purposes here, since no new messages were allowe to be fetched and there should be only one version of last state). Note that the $suspend$ flags are set to False for every processor when reconfiguration finishes and when a valid view proposal is found (line 9). □

The above lemmas lead to the following theorem.

Theorem 4.13 *Starting in an arbitrary state in an execution R of Algorithm 4.7, the algorithm simulates state machine replication preserving the virtual synchrony property, even in the case of reconfiguration when this is delicate, i.e., when initiated by $needDelicateReconf()$.*

Self-stabilizing reconfigurable emulation of shared memory. Birman et al. [5] show how a virtual synchrony solution can lead to a reconfigurable emulation of shared memory. Following this approach, and using our self-stabilizing reconfigurable SMR solution discussed in this section, and our increment counter scheme (Section 4.2), we can obtain a self-stabilizing reconfigurable emulation of shared memory. Given a conflict-free configuration, a typical two-phase read and write protocol can be used for the shared memory emulation. In the event of a delicate reconfiguration, the coordinator (of the virtual synchrony algorithm) suspends reads and writes on the register and once a new configuration is established, the emulation continues. Virtual synchrony ensures that the state of the system, in this case the state of the object, is preserved (c.f., Theorem 4.13). In the event of a brute force reconfiguration (e.g. due to transient faults or violation of the churn rate), the system will automatically recover and eventually reach a legal execution (in this case the state of the system may be lost).

We note that our proposed self-stablizing reconfigurable SMR and shared memory emulation solutions are suspending, in the sense that they do not provide service during a reconfiguration. With some extra care and under certain conditions we believe that they can be modified to provide continuous service, but in general, it remains an interesting open question whether a *self-stabilizing* service, such as reconfigurable SMR or distributed shared memory that does not suspend, is possible. In [5], Birman et al. discuss the tradeoffs of suspending and non suspending reconfiguration (such as the ones provided in [17] and [2]). It is argued, that suspending services provide some simpler solutions, and may be enhanced for more efficient reconfiguration decisions so that the time for reconfiguration and state transfer before reconfiguration can be reduced.

5 Conclusion

We presented the first self-stabilizing reconfiguration scheme that recovers automatically from transient faults, such as temporary violation of the predefined churn rate or the unexpected activities of processors and communication channels, using a bounded amount of local storage and message size. We showed

how this scheme can be used for the implementation of several dynamic distributed services, such as a self-stabilizing reconfigurable virtual synchrony, which in turn can be used for developing self-stabilizing reconfigurable SMR and shared memory emulation solutions. We use a number of bootstrapping techniques for allowing the system to always recover from arbitrary transient faults, for example, when the current configuration includes no active processors. We believe that the presented techniques provide a generic blueprint for different solutions that are needed in the area of self-stabilizing high-level communication and synchronization primitives, which need to deal with processor joins and leaves as well as transient faults.

References

[1] Marcos K. Aguilera, Idit Keidar, Dahlia Malkhi, Jean-Philippe Martin, and Alexander Shraer. Reconfiguring replicated atomic storage: A tutorial. *Bulletin of the EATCS*, 102:84–108, 2010.

[2] Marcos Kawazoe Aguilera, Idit Keidar, Dahlia Malkhi, and Alexander Shraer. Dynamic atomic storage without consensus. *J. ACM*, 58(2):7, 2011.

[3] Hagit Attiya, Hyun Chul Chung, Faith Ellen, Saptaparni Kumar, and Jennifer L. Welch. Simulating a shared register in an asynchronous system that never stops changing - (extended abstract). In *Distributed Computing - 29th International Symposium, DISC 2015, Tokyo, Japan, October 7-9, 2015, Proceedings*, pages 75–91, 2015.

[4] R. Baldoni, S. Bonomi, A. M. Kermarrec, and M. Raynal. Implementing a register in a dynamic distributed system. In *Distributed Computing Systems, 2009. ICDCS '09. 29th IEEE International Conference on*, pages 639–647, June 2009.

[5] Ken Birman, Dahlia Malkhi, and Robbert van Renesse. Virtually synchronous methodology for dynamic service replication. Technical Report MSR-TR-2010-151, Microsoft Research, 2010.

[6] Peva Blanchard, Shlomi Dolev, Joffroy Beauquier, and Sylvie Delaët. Practically self-stabilizing paxos replicated state-machine. In *In Revised Selected Papers of the Second International Conference on Networked Systems, NETYS 2014*, volume 8593 of *Lecture Notes in Computer Science*, pages 99–121. Springer, 2014.

[7] Vita Bortnikov, Gregory Chockler, Dmitri Perelman, Alexey Roytman, Shlomit Shachor, and Ilya Shnayderman. Reconfigurable state machine replication from non-reconfigurable building blocks. *CoRR*, abs/1512.08943, 2015.

[8] Gregory V. Chockler, Seth Gilbert, Vincent Gramoli, Peter M. Musial, and Alexander A. Shvartsman. Reconfigurable distributed storage for dynamic networks. *J. Parallel Distrib. Comput.*, 69(1):100–116, 2009.

[9] Shlomi Dolev. *Self-stabilization*. The MIT press, 2000.

[10] Shlomi Dolev, Swan Dubois, Maria Potop-Butucaru, and Sébastien Tixeuil. Stabilizing data-link over non-fifo channels with optimal fault-resilience. *Inf. Process. Lett.*, 111(18):912–920, 2011.

[11] Shlomi Dolev, Chryssis Georgiou, Ioannis Marcoullis, and Elad Michael Schiller. Self-stabilizing virtual synchrony. In *Proceedings of the 17th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS 2015*, pages 248–264, 2015.

[12] Shlomi Dolev, Ariel Hanemann, Elad Michael Schiller, and Shantanu Sharma. Self-stabilizing end-to-end communication in (bounded capacity, omitting, duplicating and non-fifo) dynamic networks - (extended abstract). In *Proceedings of the 14th International Symposium Stabilization, Safety, and Security of Distributed Systems, SSS 2012*, pages 133–147, 2012.

[13] Shlomi Dolev and Ted Herman. Superstabilizing protocols for dynamic distributed systems. *Chicago J. Theor. Comput. Sci.*, 1997, 1997.

[14] Shlomi Dolev, Elad Schiller, and Jennifer L. Welch. Random walk for self-stabilizing group communication in ad hoc networks. *IEEE Trans. Mob. Comput.*, 5(7):893–905, 2006.

[15] Shlomi Dolev and Nir Tzachar. Empire of colonies: Self-stabilizing and self-organizing distributed algorithm. *Theor. Comput. Sci.*, 410(6-7):514–532, 2009.

[16] Eli Gafni and Dahlia Malkhi. Elastic configuration maintenance via a parsimonious speculating snapshot solution. In *Distributed Computing - 29th International Symposium, DISC 2015, Tokyo, Japan, October 7-9, 2015, Proceedings*, pages 140–153, 2015.

[17] Seth Gilbert, Nancy A. Lynch, and Alexander A. Shvartsman. Rambo: a robust, reconfigurable atomic memory service for dynamic networks. *Distributed Computing*, 23(4):225–272, 2010.

[18] Leander Jehl, Roman Vitenberg, and Hein Meling. Smartmerge: A new approach to reconfiguration for atomic storage. In *Distributed Computing - 29th International Symposium, DISC 2015, Tokyo, Japan, October 7-9, 2015, Proceedings*, pages 154–169, 2015.

[19] Leslie Lamport, Dahlia Malkhi, and Lidong Zhou. Reconfiguring a state machine. *SIGACT News*, 41(1):63–73, 2010.

[20] Peter M. Musial, Nicolas C. Nicolaou, and Alexander A. Shvartsman. Implementing distributed shared memory for dynamic networks. *Commun. ACM*, 57(6):88–98, 2014.

[21] David Peleg and Avishai Wool. Crumbling walls: A class of practical and efficient quorum systems. *Distributed Computing*, 10(2):87–97, 1997.

[22] Alexander Spiegelman, Idit Keidar, and Dahlia Malkhi. Dynamic reconfiguration: A tutorial. *OPODIS 2015*, 2015.

[23] Marko Vukolic. *Quorum Systems: With Applications to Storage and Consensus*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2012.